EXAMTOPICS

-Expert Verified, Online, **Free**.



CERTIFICATION TEST

-CertificationTest.net - Cheap & Quality Resources With Best Support

Question #1 *Topic 1*

Which three are true about thresholds, metrics, and server-generated alerts? (Choose three.)

A. All metrics are instance related.

B. Cleared stateful alerts are displayed by querying DBA_ALERT_HISTORY.

C. A space usage management alert is automatically cleared after the underlying problem is resolved.

D. They are generated by SMON when a tablespace is 97% full.

E. Metrics are statistical counts for a specific unit.

F. STATISTICS_LEVEL must be set to ALL to generate alerts.

Suggested Answer: *C, E, F*
Reference: https://docs.oracle.com/cd/E11882_01/server.112/e41573/autostat.htm#PFGRF027 https://blogs.oracle.com/oem/how-to-clear-an-alert-v2

Community vote distribution
BCE (100%)

janw Highly Voted 5 years, 2 months ago
correct: BCE
upvoted 16 times

Alejandrrro Highly Voted 4 years, 7 months ago
Correct: B,C,E

Reference: Oracle Database 12c: Administration Wrokshop

A. False. Tablespace Space Usage metris is database related B. True. Stateful alerts appear in DBA_OUTSTANDING_ALERTS and, when cleared, go to DBA_ALERT_HISTORY C. True. Most alerts (such as "Out of space") are clared automatically when the cause of the problem disappears. D. False. They're generated by MMON E. True F. False. This parameter is not required to generate alerts
upvoted 9 times

gio Most Recent 2 years, 1 month ago
Selected Answer: BCE
Sure of BCE
upvoted 2 times

ScottL 2 years, 3 months ago
Selected Answer: BCE
BCE Correct
upvoted 1 times

ScottL 2 years, 4 months ago
Selected Answer: BCE
Correct BCE
upvoted 1 times

raferen10 2 years, 5 months ago
Selected Answer: BCE
Correct: B,C,E A. False. Tablespace Space Usage metris is database related B. True. Stateful alerts appear in DBA_OUTSTANDING_ALERTS and, when cleared, go to DBA_ALERT_HISTORY C. True. Most alerts (such as "Out of space") are clared automatically when the cause of the problem disappears. D. False. generated by MMON E. True. Its definition of statistics F. False. This parameter is not required to generate alerts
upvoted 1 times

vkra 2 years, 6 months ago
Selected Answer: BCE
BCE is correct
upvoted 1 times

hilaire 2 years, 8 months ago

Correct BCE

upvoted 1 times

MKKB 2 years, 12 months ago

Selected Answer: BCE

Correct BCE

upvoted 1 times

chefdba 3 years, 10 months ago

BCE the right answer

upvoted 1 times

Neil107 4 years ago

Correct

upvoted 1 times

SMagny 4 years, 10 months ago

F is incorrect. STATISTICS_LEVEL=TYPICAL is sufficient. Answer is B,C,E.

upvoted 1 times

cerebro2000x 5 years ago

STATISTICAL_LEVEL=TYPICAL is ok

upvoted 2 times

Jatindra 5 years, 2 months ago

Reference Student guide: Threshold alerts are also referred to as stateful alerts which are automatically cleared when an alert condition clears. Stateful alerts appear in DBA_OUTSTANDING_ALERTS and when cleared go to DBA_ALERT_HISTORY.

upvoted 2 times

janw 5 years, 2 months ago

F incorrect ==> typical is sufficient

upvoted 4 times

janw 5 years, 2 months ago

C correct

Most alerts, such as the Out of Space alert, are cleared automatically when the cause of the problem disappears.

upvoted 1 times

janw 5 years, 2 months ago

B correct DBA_ALERT_HISTORY describes a time-limited history of alerts which are no longer outstanding.

upvoted 1 times

Question #2 *Topic 1*

While backing up to an SBT channel, you determine that the read phase of your compressed Recovery Manager (RMAN) incremental level 0 backup is a bottleneck.
FORCE LOGGING is enabled for the database.
Which two could improve read performance? (Choose two.)

A. Increase the size of tape I/O buffers.

B. Disable FORCE LOGGING for the database.

C. Increase the size of the database buffer cache.

D. Enable asynchronous disk I/O.

E. Increase the level of RMAN multiplexing.

Suggested Answer: *C, E*
Official Answer: CE

janw Highly Voted 6 years, 2 months ago
DE correct C = incorrect, rman does not use database buffer cache
upvoted 14 times

cerebro2000x 6 years ago
AE, io on tape is sync
upvoted 1 times

janw 6 years ago
A) is used for copy/write tuning

https://docs.oracle.com/html/E10642_06/rcmtunin.htm#i1006396 tape is sync, but read is on disks, so async IO can be used for reading blocks
upvoted 1 times

ama 5 years, 11 months ago
I agree, D is correct >

In RMAN operation, asynchronous Disk I/O is always better than synchronous I/O as a server process can begin an I/O and then perform other work while waiting for the I/O to complete. On operating systems that do not support native asynchronous I/O, the database can simulate it with special I/O slave processes (DBWR_IO_SLAVES)
upvoted 2 times

Alejandrrro 5 years, 7 months ago
Agree that rman doesn't use db buffer cache:

RMAN Backup Performance (Doc ID 360443.1) "RMAN uses its own buffer cache separate from the database buffer cache"
upvoted 3 times

bh19z1bt Most Recent 1 year, 10 months ago
A) is used for copy/write tuning

tape is sync, but read is on disks, so async IO can be used for reading blocks
upvoted 1 times

raferen10 1 year, 10 months ago
Selected Answer: DE
Correct: D,E

A. False. SBT buffer can't help on READ phase B. False. Force logging causes all operations in the database to be written to files, but it has no relation to RMAN reading performance C. False. RMAN uses its own buffer cache separate from the database buffer cache D. True. In RMAN operation, asynchronous Disk I/O is always better than synchronous I/O as a server process can begin an I/O and then perform other work while waiting for the

I/O to complete.On operating systems that do not support native asynchronous I/O, the database can simulate it with special I/O slave processes (DBWR_IO_SLAVES) E. True. Parallel read

upvoted 1 times

gio 3 years, 1 month ago

Selected Answer: DE

DE 100% sure

upvoted 1 times

ScottL 3 years, 3 months ago

Selected Answer: DE

DE - see Raferen10

upvoted 1 times

Jeremy81 3 years, 9 months ago

how do we enable async io?

upvoted 1 times

JESUSBB 3 years, 10 months ago

D and E are correct

upvoted 1 times

suend 4 years, 4 months ago

Selected Answer: DE

RMAN is not using db buffere so C is incorrect

upvoted 3 times

Neil107 5 years ago

It's DE.

upvoted 2 times

Dhanushbh 5 years, 1 month ago

DE correct https://docs.oracle.com/en/database/oracle/oracle-database/19/bradv/tuning-rman-performance.html#GUID-3303B929-0747-4E78-BED1-9F468213F357

upvoted 1 times

SMagny 5 years, 10 months ago

C is incorrect. RMAN does not use database buffer cache. The answer is D,E. Asynchronous I/O (D) and multiplexing (Parallel read - E) improve read performance.

upvoted 3 times

yurijk 5 years, 11 months ago

AE. D async is for writes, not reads.

upvoted 1 times

ama 5 years, 11 months ago

No, you are totally wrong…

Synchronous and Asynchronous Disk I/O When a channel reads from or writes to disk, the I/O is either synchronous I/O or asynchronous I/O. When the disk I/O is synchronous, a server process can perform only one task at a time. When the disk I/O is asynchronous, a server process can begin an I/O operation and then perform other work while waiting for the I/O to complete. RMAN can also begin multiple I/O operations before waiting for the first to complete.

Take a look at https://docs.oracle.com/database/121/BRADV/rcmtunin.htm#BRADV90067

and Focus on the Read-Phase

upvoted 4 times

yurijk 5 years, 11 months ago

Yep, you're right. async can help at READ phase, and A. SBT buffer can't help on READ phase. So finally right answer are: DE

upvoted 3 times

janw 6 years, 2 months ago

When a channel reads from or writes to disk, the I/O is either synchronous I/O or asynchronous I/O. When the disk I/O is synchronous, a server process can perform only one task at a time. When the disk I/O is asynchronous, a server process can begin an I/O operation and then perform other work while waiting for the I/O to complete. RMAN can also begin multiple I/O operations before waiting for the first to complete.

When reading from an ASM disk group, you should use asynchronous disk I/O if possible. Also, if a channel reads from a raw device managed with a volume manager, then asynchronous disk I/O also works well. Some operating systems support native asynchronous disk I/O. The database takes advantage of this feature if it is available.

upvoted 4 times

janw 6 years, 2 months ago

You can improve performance by increasing the degree of multiplexing used for backing up. This increases the rate at which RMAN fills tape buffers, which makes it more likely that buffers are sent to the media manager fast enough to maintain streaming.

upvoted 3 times

janw 6 years, 2 months ago

CE correct

upvoted 2 times

xRodge 5 years, 1 month ago

Do you mean to say (based on your analysis), that DE is correct?

D. Enable asynchronous disk I/O. E. Increase the level of RMAN multiplexing.

upvoted 1 times

Question #3 *Topic 1*

For which two requirements can you use the USER_TABLESPACE clause with the CREATE PLUGGABLE DATABASE command? (Choose two.)

A. to specify a default tablespace in a PDB cloned from another PDB in the same CDB.

B. to exclude all tablespaces except SYSTEM, SYSAUX, and TEMP when plugging in a PDB

C. to include specific user tablespaces only when relocating a PDB

D. to specify the list of user tablespaces to include when moving a non-CDB to a PDB

E. to exclude a temp tablespace when plugging in a PDB

F. to specify the list of tablespaces to include when creating a PDB from the CDB seed

Suggested Answer: *E, F*

Reference: https://docs.oracle.com/en/database/oracle/oracle-database/12.2/admin/creating-and-removing-pdbs-with-sql-plus.html#GUID-1C47D543-8376-48AE- A1AE-632316731D59

Community vote distribution

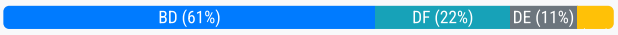

janw Highly Voted 6 years, 2 months ago

B. to exclude all tablespaces except SYSTEM, SYSAUX, and TEMP when plugging in a PDB D. to specify the list of user tablespaces to include when moving a non-CDB to a PDB

upvoted 21 times

pedro0986 6 years, 2 months ago

It's correct BD?

upvoted 1 times

Alejandrrro 5 years, 7 months ago

To confirm D tried plugging in a PDB in my test system with USER_TABLESPACES clause and an excluded tablespace is indeed offline.

CREATE PLUGGABLE DATABASE XEPDB2 USING '/opt/oracle/oradata/XE/XEPDB2/XEPDB2.xml' USER_TABLESPACES=NONE NOCOPY TEMPFILE REUSE;

SQL> select TABLESPACE_NAME,status from dba_tablespaces;

TABLESPACE_NAME STATUS ---------- -------- SYSTEM ONLINE SYSAUX ONLINE UNDOTBS1 ONLINE TEMP ONLINE CUSTOM OFFLINE

upvoted 1 times

ama Highly Voted 5 years, 6 months ago

i double checked B again.

B. to exclude all tablespaces except SYSTEM, SYSAUX, and TEMP when plugging in a PDB ?

Its definitely FALSE

Because to exclude some tablespaces we can use for example USER_TABLESPACES=ALL EXCEPT('tbs1','tbs4','tbs5');

in this case all tablespaces will be included except tbs1, tbs4 and tbs5

Howver This clause does not apply to the SYSTEM, SYSAUX, or TEMP tablespaces. so you should not include these tablespaces in a tablespace list for this clause.

If i want to translate Answer B into an example, it should look like that:

USER_TABLESPACES=ALL EXCEPT('system','sysaux','emp');

Apparently something like that is not allowed, therefor Answer B is wrong

Correct answers are D, F

upvoted 6 times

raferen10 3 years, 5 months ago

if the NONE option is specified, all tablespaces except SYSTEM, SYSAUX, or TEMP are excluded since this clause does not apply to these tablespaces

upvoted 3 times

sankaran1 Most Recent 1 year, 4 months ago

Selected Answer: BD

B: please read the keyword "except" , that makes this statement true. Reference : https://docs.oracle.com/en/database/oracle/oracle-database/19/sqlrf/CREATE-PLUGGABLE-DATABASE.html

upvoted 1 times

cujar2003 1 year, 10 months ago

Selected Answer: BD

https://oracle-base.com/articles/12c/multitenant-user_tablespaces-clause-12c

upvoted 1 times

frankzeng 3 years, 2 months ago

from https://docs.oracle.com/en/database/oracle/oracle-database/19/sqlrf/CREATE-PLUGGABLE-DATABASE.html#GUID-F2DBA8DD-EEA8-4BB7-A07F-78DC04DB1FFC, there is no user_tablespaces in syntax. F is wrong

upvoted 1 times

ScottL 3 years, 3 months ago

Selected Answer: BD

B and D correct.

F incorrect based on my reading of the doc - "This clause lets you specify the tablespaces to be made available in the new PDB. The SYSTEM, SYSAUX, and TEMP tablespaces are available in all PDBs and cannot be specified in this clause. "

https://docs.oracle.com/en/database/oracle/oracle-database/19/sqlrf/CREATE-PLUGGABLE-DATABASE.html#GUID-F2DBA8DD-EEA8-4BB7-A07F-78DC04DB1FFC

upvoted 2 times

gio 3 years, 4 months ago

Selected Answer: BD

as explained here: https://docs.oracle.com/en/database/oracle/oracle-database/19/sqlrf/CREATE-PLUGGABLE-DATABASE.html#GUID-F2DBA8DD-EEA8-4BB7-A07F-78DC04DB1FFC

upvoted 1 times

ErikJanssen 3 years, 5 months ago

I think BD as well. I looked in the syntax diagrams and the user_tablespace clause is also valid for answer F.
https://docs.oracle.com/en/database/oracle/oracle-database/19/sqlrf/CREATE-PLUGGABLE-DATABASE.html#GUID-F2DBA8DD-EEA8-4BB7-A07F-78DC04DB1FFC

upvoted 1 times

jackymak 3 years, 5 months ago

Selected Answer: BF

https://oracle-base.com/articles/12c/multitenant-user_tablespaces-clause-12c Remember, the SYSTEM, SYSAUX and TEMP tablespaces are always included.

upvoted 1 times

raferen10 3 years, 5 months ago

Selected Answer: DE

Correct: B,D

A. False. Default tablespaces are defined with the DEFAULT_TABLESPACE clause not with USER_TABLESPACE clause. B. True. Specify NONE in USER_TABLESPACE clause to make only the SYSTEM, SYSAUX, and TEMP tablespaces available in the new PDB. C. False. The USER_TABLESPACE clause is available for use in 3 scenarios: When creating a PDB from SEED, When cloning an existing PDB or When creating a PDB from XML File D. True. When you move a non-CDB to a PDB, and the non-CDB had several schemas that each supported a different application, you can use this clause

to separate the data belonging to each schema into a separate PDB. This technique assumes that each schema used a separate tablespace in the non-CDB. E. False. The SYSTEM, SYSAUX, and TEMP tablespaces always are available in all PDBs and cannot be specified in this clause F. False. When create PDB from seed no users tablespaces are defined in seed container

upvoted 2 times

psebik 3 years, 7 months ago

Hello. Answer F is false. Database cdb@seed have only dictionary structures. This database have only system tablespaces (SYSTEM, SYSAUX ...). In this time we omit user_tablespace clause. https://docs.oracle.com/database/121/ADMIN/cdb_plug.htm#ADMIN13850 This is tricky question. :). Corect are B and D.

upvoted 2 times

Yennism 3 years, 9 months ago

Selected Answer: BD

B= with NONE D=Specify tablespace to make the tablespace available in the new PDB

upvoted 1 times

ErikJanssen 3 years, 10 months ago

I found https://docs.oracle.com/en/database/oracle/oracle-database/19/sqlrf/CREATE-PLUGGABLE-DATABASE.html#GUID-F2DBA8DD-EEA8-4BB7-A07F-78DC04DB1FFC very helpfull. See user_tablespaces_clause

upvoted 1 times

vkra 3 years, 5 months ago

Not C : USER_TABLESPACE is a clause for CREATE PLUGGABLE DATABASE options: (create_pdb_from_seed, create_pdb_clone, create_pdb_from_xml) and not an option for RELOCATION clause.

upvoted 1 times

Franky_T 4 years, 3 months ago

This one is tricky, and I found 3 correct answers. The USER_TABLESPACE clause is available for use in 3 scenarios. A. When creating a PDB from SEED. B. When cloning an existing PDB. C. When creating a PDB from XML File USER_TABLESPACE has 4 choices: (This clause lets you specify the tablespaces to be made available in the new PDB. The SYSTEM, SYSAUX, and TEMP tablespaces are available in all PDBs and cannot be specified in this clause.) 1. Specify tablespace to make the tablespace available in the new PDB. You can specify more than one tablespace in a comma-separated list. 2. Specify ALL to make all tablespaces available in the new PDB. This is the default. 3. Specify ALL EXCEPT to make all tablespaces available in the new PDB, except the specified tablespaces. 4. Specify NONE to make only the SYSTEM, SYSAUX, and TEMP tablespaces available in the new PDB.

A. is wrong. Default tablespaces are defined with the DEFAULT_TABLESPACE clause. B is correct. See point 4 for USER_TABLESPACE. C is wrong. See description before point 1. D is correct. See point 1. E is wrong. See description before point 1. F is correct. See point 1.

upvoted 3 times

jackymak 4 years, 1 month ago

Will B be incorrect? Cause it's 'when plugging in a PDB', Not cloning or creating.

upvoted 1 times

Datajimm 4 years, 4 months ago

Selected Answer: BD

I'd say B is correct. The documentation says: "Specify NONE to exclude all tablespaces." and later "This clause does not apply to the SYSTEM, SYSAUX, or TEMP tablespaces." That is, when specifying NONE, only SYSTEM, SYSAUX and TEMP will be included. And that's what answer B says.

https://docs.oracle.com/en/database/oracle/oracle-database/19/multi/overview-of-pdb-creation.html#GUID-1C47D543-8376-48AE-A1AE-632316731D59

upvoted 5 times

FeaRoX 4 years, 4 months ago

assuming it should be "PDB seed" in F :

SQL> create pluggable database pdb4 admin user pdbadmin identified by pdbadmin USER_TABLESPACES=('USERS');

Pluggable database created.

Not much sense, but it works.

upvoted 3 times

casfdsaf 4 years, 4 months ago

Selected Answer: DF

DF i think

upvoted 2 times

Question #4 *Topic 1*

Which three are true about requirements for various FLASHBACK operations? (Choose three.)

A. FLASHBACK transaction query requires undo to retrieve all versions of a row that existed between two points in time.

B. FLASHBACK drop requires that the RECYCLEBIN parameter be set to ON.

C. FLASHBACK version query requires that the RECYCLEBIN parameter be set to ON.

D. FLASHBACK DATA ARCHIVE requires undo to store all versions of all rows of a table being tracked.

E. FLASHBACK drop requires undo to retrieve all versions of a row that existed between two points in time.

F. FLASHBACK version query requires undo to retrieve all versions of a row that existed between two points in time.

Suggested Answer: *A, B, C*
Reference: https://books.google.com.pk/books?id=0iwrL9P25Z0C&pg=PA35&lpg=PA35&dq=FLASHBACK+transaction+query+requires+undo+to+retrieve+all+versions+of+a+row+that+existed+between+two+points+in+time&source=bl&ots=MJnYl5CZ1u&sig=ACfU3U0dWP-NPd-U8uu3zbaoi3YZzT0FTQ&hl=en&sa=X&ved=2ahUKEwigi5_Jl9joAhXzUBUIHUh4DksQ6AEwAnoECAsQJg#v=onepage&q=FLASHBACK%20transaction%20query%20requires%20undo%20to%20retrieve%20all%20versions%20of%20a%20row%20that%20existed%20between%20two%20points%20in%20time&f=false https://docs.oracle.com/cd/E18283_01/server.112/e17120/tables011.htm

Community vote distribution

BDF (54%) ABF (46%)

janw Highly Voted 6 years, 2 months ago
A B F

C incorrect, recycle bin not used for flashback version
upvoted 23 times

ama 5 years, 11 months ago
i agree!!!
upvoted 3 times

yarsalan Highly Voted 4 years, 6 months ago
Selected Answer: BDF
In my opinion, B,D,F are correct
upvoted 6 times

sankaran1 Most Recent 1 year, 4 months ago
Selected Answer: BDF
Why A is wrong because it will not use two point in time. It will use only one given time https://docs.oracle.com/en/database/oracle/oracle-database/19/adfns/flashback.html#GUID-212A4073-ABE2-4053-A852-3A41DEF6A24D
upvoted 1 times

219f72b 1 year, 5 months ago
Selected Answer: BDF
Flashback transaction query can be used to get extra information about the transactions listed by flashback version queries. Flashback archive provides the ability to automatically track and archive transactional data changes to specified database objects.
upvoted 1 times

bob0069 1 year, 5 months ago
Selected Answer: ABF
D - does not need undo
upvoted 1 times

gauloises22 1 year, 7 months ago
Selected Answer: BDF

Last, UNDO is used for all Flashback features: "Oracle Flashback features use the Automatic Undo Management (AUM) system to obtain metadata and historical data for transactions. They rely on undo data, which are records of the effects of individual transactions. For example, if a user runs an UPDATE statement to change a salary from 1000 to 1100, then Oracle Database stores the value 1000 in the undo data. Undo data is persistent and survives a database shutdown. It is retained for the time specified by undo_retention, or up to the tuned undo retention in the presence of Automatic Undo Management (AUM). By using flashback features, you can use undo data to query past data or recover from logical damage."

upvoted 1 times

gauloises22 1 year, 7 months ago

Selected Answer: BDF

D is correct because the Time Travel (Archiver) is using tablespaces, but it is enabled per table: "Flashback Time Travel provides the ability to track and store transactional changes to a table over its lifetime... A Flashback Archive consists of one or more tablespaces or parts thereof... By default, Flashback Archive is not enabled for any tables. You can enable Flashback Archive for a table if all of these conditions are true:... When choosing a Flashback Archive for a specific table, consider the data retention requirements for the table and the retention times of the Flashback Archives on which you have the FLASHBACK ARCHIVE object privilege."

upvoted 1 times

gauloises22 1 year, 7 months ago

Selected Answer: BDF

It is not A because flashback transaction query don't use two points in time necessarily, it is just a query for transaction metadata. You can utilize Flashback Version Query ("VERSIONS BETWEEN TIMESTAMP") as a subquery in a Flashback Transaction Query to associate each row version with the LOGON_USER responsible for the row data change "Oracle Flashback Transaction Query: Use this feature to retrieve metadata and historical data for a given transaction or for all transactions in a given time interval. To perform an Oracle Flashback Transaction Query, select from the static data dictionary view FLASHBACK_TRANSACTION_QUERY."

upvoted 1 times

pitt84 1 year, 10 months ago

BDF

upvoted 1 times

antonica 2 years, 4 months ago

ABF, I AGREE

upvoted 1 times

dilmaj 2 years, 7 months ago

ABF is correct. D is not correct because it's about flashback of a Tablespace not table: A flashback data archive consists of one or more tablespaces or parts of tablespaces. When you create a flashback data archive, you specify the name, retention period, and tablespace

upvoted 1 times

guretto 2 years, 7 months ago

ABF

A true --> FTQ uses UNDO. B true --> FD uses recyclebin C false --> FVQ does not use recyclebin. See F D false --> FDA use Flashback Archive Tablespace E false --> FD does not use undo. See B F true --> FVQ uses UNDO

upvoted 1 times

Yennism 2 years, 8 months ago

I read the note Click to add to Favorites FDA - Flashback Data Archive Usage and Best Practices (a.k.a. Total Recall) (Doc ID 2370465.1). and use UNDO, then the aswer is BDF

upvoted 2 times

auwia 2 years, 10 months ago

Selected Answer: ABF

A, B, and F are correct. C is wrong because FDA needs a user tablespace.

upvoted 1 times

gio 3 years, 1 month ago

Selected Answer: ABF

ABF 100 %

upvoted 2 times

kaka321 3 years, 1 month ago

ABF is correct. A which is the most debated is right I just confirm from oracle document that flashback transaction query uses time interval

upvoted 1 times

Guhborges 3 years, 1 month ago

Selected Answer: BDF

Why D? : Oracle Flashback features uses the Automatic Undo Management (AUM) system to obtain metadata and historical data for transactions. They rely on undo data, which are records of the effects of individual transactions.

upvoted 1 times

Question #5 *Topic 1*

Which three actions are performed by the Oracle Preinstallation RPM, oracle-database-server-xxxx-preinstall, for Oracle Grid Infrastructure, where xxxx is the Oracle version and release? (Choose three.)

A. performing checks to ensure minimum configuration requirements for Oracle Grid Infrastructure are met

B. creating the oracle OS user

C. creating the OSDBA (dba) group

D. creating the oraInventory (oinstall) group

E. creating the grid OS user

F. configuring the OS for Oracle Automatic Storage Management shared storage access

Suggested Answer: *C, D, F*

Reference: https://docs.oracle.com/cd/E11882_01/install.112/e41961/prelinux.htm#CWLIN2932

Community vote distribution

BCD (100%)

janw Highly Voted 5 years, 8 months ago

correct = BCD

upvoted 20 times

asefa Highly Voted 5 years, 8 months ago

janW You are right the answer is BCD.

upvoted 6 times

antonica Most Recent 1 year, 11 months ago

BCD, Automatically downloads and installs any additional RPM packages needed for installing Oracle Grid Infrastructure and Oracle Database, and resolves any dependencies. Creates an oracle user, and creates the oraInventory (oinstall) and OSDBA (dba) groups for that user. As needed, sets sysctl.

https://docs.oracle.com/en/database/oracle/oracle-database/18/cwlin/about-the-oracle-preinstallation-rpm.html#GUID-C15A642B-534D-4E4A-BDE8-6DC7772AA9C8

upvoted 1 times

3c4de23 2 years, 1 month ago

I am really disappointed by the answer given for this unless the question mixed grid and database install.. the correct answer is Oracle user, OS Group, and OraInventory

upvoted 1 times

gio 2 years, 7 months ago

Selected Answer: BCD

100% BCD

upvoted 2 times

raferen10 2 years, 11 months ago

Selected Answer: BCD

Correct: B,C,D When installed, the Oracle Database Preinstallation RPM does the following: -Automatically downloads and installs any additional RPM packages needed for installing Oracle Grid Infrastructure and Oracle Database, and resolves any dependencies -Creates an oracle user, and creates the oraInventory (oinstall) and OSDBA (dba) groups for that user -As needed, sets sysctl.conf settings, system startup parameters, and driver parameters to values based on recommendations from the Oracle Database Preinstallation RPM program -Sets hard and soft resource limits -Sets other recommended parameters, depending on your kernel version -Sets numa=off in the kernel for Linux x86_64 machines.
https://docs.oracle.com/en/database/oracle/oracle-database/19/ladbi/about-the-oracle-preinstallation-rpm.html

upvoted 4 times

jareach 3 years ago

Selected Answer: BCD

The RPM does not configure the system

upvoted 1 times

DonqpeR 3 years, 1 month ago

Selected Answer: BCD

https://docs.oracle.com/en/database/oracle/oracle-database/19/ladbi/about-the-oracle-preinstallation-rpm.html

upvoted 1 times

477267 3 years, 1 month ago

Correct BCD, https://docs.oracle.com/en/database/oracle/oracle-database/18/ladbi/about-the-oracle-preinstallation-rpm.html#GUID-C15A642B-534D-4E4A-BDE8-6DC7772AA9C8

upvoted 1 times

477267 3 years, 1 month ago

https://docs.oracle.com/en/database/oracle/oracle-database/18/ladbi/about-the-oracle-preinstallation-rpm.html#GUID-C15A642B-534D-4E4A-BDE8-6DC7772AA9C8

upvoted 1 times

Patrick9230 3 years, 7 months ago

Selected Answer: BCD

BCD is correct

upvoted 1 times

dasigrist 4 years, 10 months ago

I agree that B, C, D are correct, but what is wrong with A?

upvoted 2 times

asduta 3 years, 8 months ago

I think that one will be run during the Grid installation prereq check.

upvoted 1 times

Bda 4 years, 11 months ago

Pls help.

For B, what is meant by 'creating "the" oracle OS user?'

I have doubt why B is correct and F is incorrect.

https://docs.oracle.com/en/database/oracle/oracle-database/12.2/ladbi/overview-of-oracle-linux-configuration-with-oracle-rpms.html#GUID-693599D4-BD32-4E6A-9689-FA7D1CD75653

Configuring a server using Oracle Linux and the Oracle Preinstallation RPM consists of the following steps:

1. Install Oracle Linux.

2. Register your Linux distribution with Oracle Unbreakable Linux Network (ULN) or download and configure the yum repository for your system using the Oracle Linux yum server for your Oracle Linux release.

3. Install the Oracle Preinstallation RPM with the RPM for your Oracle Grid Infrastructure and Oracle Database releases, and update your Linux release.

4. Create role-allocated groups and users with identical names and ID numbers.

5. Complete network interface configuration for each cluster node candidate.

6. Complete system configuration for shared storage access as required for each standard or Core node cluster candidate.

upvoted 1 times

SMagny 5 years, 4 months ago

The answer is B,D,C. F is incorrect.

upvoted 2 times

IOracle 5 years, 7 months ago

B,C,D - from here https://docs.oracle.com/cd/E11882_01/install.112/e41961/prelinux.htm#CWLIN172

upvoted 4 times

janw 5 years, 8 months ago

When installed, the Oracle Preinstallation RPM does the following:

Automatically downloads and installs any additional RPM packages needed for installing Oracle Grid Infrastructure and Oracle Database, and resolves any dependencies

Creates an oracle user, and creates the oraInventory (oinstall) and OSDBA (dba) groups for that user

As needed, sets sysctl.conf settings, system startup parameters, and driver parameters to values based on recommendations from the Oracle Preinstallation RPM program

Sets hard and soft resource limits

Sets other recommended parameters, depending on your kernel version

Sets numa=off in the kernel for Linux x86_64 machines.

upvoted 6 times

Question #6 *Topic 1*

Which two are true about common objects? (Choose two.)

A. They can be created only in CDB$ROOT.

B. They can be only metadata-linked in an application container.

C. They can exist in user-defined schemas only in application containers.

D. They can exist in CDB$ROOT and an application root.

E. They can be extended data-linked in CDB$ROOT.

F. They can be created only in an application root.

Suggested Answer: *E, F*

Reference: https://blog.toadworld.com/2017/08/01/oracle-multi-tenant-application-containers-part-iii-sharing-of-data-in-application-common-objects

Community vote distribution

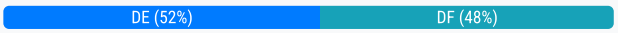

erial Highly Voted 5 years, 11 months ago

READ THIS: A common object is defined in either the CDB root or an application root, and can be referenced using metadata links or object links. A local object is every object that is not a common object.

Database-supplied common objects are defined in CDB$ROOT and cannot be changed. Oracle Database does not support creation of common objects in CDB$ROOT.

You can create most schema objects—such as tables, views, PL/SQL and Java program units, sequences, and so on—as common objects in an application root. If the object exists in an application root, then it is called an application common object. ACCORDING THE LECTURE, CORRECT ANSWERS ARE: DF D: they can exist in CDB$ROOT and application root. F: they can be created only in application root.

upvoted 15 times

jareach 3 years, 6 months ago

However F contradicts D...

upvoted 1 times

ErikJanssen 3 years, 5 months ago

I think D and F are the correct answers. They do not contradict; You must distinguish here between Oracle supplied objects already present in the root after install and the common objects the user creates. So answer D is ok for already existing objects. And answer F is ok because users are only allowed to create common objects in application root.

upvoted 2 times

Jatindra Highly Voted 6 years, 2 months ago

Correct answer is C,D

upvoted 9 times

sankaran1 Most Recent 1 year, 4 months ago

Selected Answer: DE

F is wrong since its says "only"

upvoted 2 times

Swall 1 year, 6 months ago

Selected Answer: DF

Overview of Common and Local Objects in a CDB A common object is defined in either the CDB root or an application root, and can be referenced using metadata links or object links. A local object is every object that is not a common object.

Database-supplied common objects are defined in CDB$ROOT and cannot be changed. Oracle Database does not support creation of common objects in CDB$ROOT.

https://docs.oracle.com/en/database/oracle/oracle-database/19/multi/overview-of-the-multitenant-architecture.html#GUID-8B4B4F8B-1C2B-42C3-BE77-E9AA25CA397A

upvoted 1 times

auwia 2 years, 1 month ago

Selected Answer: DE

For me F is false "F. They can be created only in an application root." --> for "normal user", but not for "oracle developers", that's why they exists! :-) we should go for D, E.

upvoted 2 times

antonica 2 years, 4 months ago

common object An object that resides either in the CDB root or an application root that shares either data (a data-linked common object) or metadata (a metadata-linked common object). All common objects in the CDB root are Oracle-supplied. A common object in an application root is called an application common object.

upvoted 1 times

antonica 2 years, 4 months ago

by definition, E is wrong, just D is true

upvoted 1 times

antonica 2 years, 4 months ago

https://docs.oracle.com/en/database/oracle/oracle-database/19/multi/glossary.html

upvoted 1 times

dancymonkey 2 years, 6 months ago

DE is correct

F is incorrect - becuase not only in an application root

https://docs.oracle.com/en/database/oracle/oracle-database/23/multi/glossary.html#GUID-1C63428A-905A-4918-95AA-F3BBA3D2B825:~:text=Parent%20topic%3A%20Glossary-,common%20object,in%20an%20application%20root%20is%20called%20an%20application%20common%20object.,-Parent%20topic%3A

common object An object that resides either in the CDB root or an application root that shares either data (a data-linked common object) or metadata (a metadata-linked common object). All common objects in the CDB root are Oracle-supplied. A common object in an application root is called an application common object.

upvoted 2 times

LEOC71 2 years, 9 months ago

Selected Answer: DF

https://docs.oracle.com/en/database/oracle/oracle-database/19/multi/overview-of-the-multitenant-architecture.html#GUID-8B4B4F8B-1C2B-42C3-BE77-E9AA25CA397A "A common object is *defined* in either the CDB root or an application root" (D: TRUE) "Oracle Database does not support creation of common objects in CDB$ROOT" + "You can create most schema objects . . . as common objects in an application root" (F:TRUE)

upvoted 1 times

auwia 2 years, 11 months ago

Selected Answer: DE

F is wrong. https://docs.oracle.com/en/database/oracle/oracle-database/12.2/cncpt/overview-of-the-multitenant-architecture.html#GUID-C7F5F6BD-E905-4402-A685-D4A79CF565EA Oracle recommends that you do not add common objects to the root or modify Oracle-supplied schemas in the root. However, you can create common users and roles for database administration.

upvoted 1 times

Guhborges 3 years, 2 months ago

Selected Answer: DE

How is it possible to be D and F? D - he says it is possible to create in CDB ROOT and in the application, in F he says it is only possible in the application root, I think people are mistaken. The correct statement is D and E.

upvoted 1 times

FelipeC 3 years, 1 month ago

D - says "exist" not "create" , you cannot create common objects in CDB$ROOT but exist supplied common objects by default , So they can only be created in aplicattion root (F)

upvoted 1 times

antonica 2 years, 4 months ago

For example, AWR data resides in the root. Each PDB uses an object link to point to the AWR data in the root

upvoted 1 times

ScottL 3 years, 3 months ago

Selected Answer: DF

I believe D,F correct. erials comments make sense

upvoted 2 times

gio 3 years, 4 months ago

Selected Answer: DF

https://docs.oracle.com/en/database/oracle/oracle-database/21/multi/application-containers2.html#GUID-CA425AC1-CA1C-498C-BC09-AD246BA67346

upvoted 2 times

raferen10 3 years, 5 months ago

Selected Answer: DF

Correct: D,F

A. Flase. Oracle Database does not support creation of common objects in CDB$ROOT. Only Oracle can create common objects in root B. False. Can be linked to metadata, data and data extended in Application Containers and Data or metadata in CDB$ROOT. C. False.Can exists in the CDB$ROOT and the Application Containers(not ONLY in Aplication Containers) but can only be CREATED in Application Containers(See A), only Oracle can create common objects that EXIST in the CDB$ROOT D. True. A common object is defined in either the CDB$ROOT(by oracle) or an application root(by users). E. False. Can be extended data liked only in application root, in CDB$root can be metadata linked and data(object) linked F. True. Database-supplied common objects are defined in CDB$ROOT and cannot be changed.

upvoted 3 times

vkra 3 years, 6 months ago

Selected Answer: DE

D eliminates A an F because of "only" and E uses data-links to make it useable

upvoted 1 times

jareach 3 years, 6 months ago

Selected Answer: DE

CDB root or application root, and can be data linked

upvoted 2 times

jareach 3 years, 6 months ago

Selected Answer: DE

F excludes D

upvoted 2 times

casfdsaf 4 years, 4 months ago

Selected Answer: DF

df i think

upvoted 1 times

Question #7 *Topic 1*

Which two are true about the Automatic Database Diagnostic Monitor (ADDM)? (Choose two.)

A. It analyzes a period of time corresponding to the 12 hours of activity.

B. It runs automatically after each AWR snapshot.

C. A DBA can run it manually.

D. Results are written to the alert log.

E. It analyzes a period of time corresponding to the last day of activity.

Suggested Answer: *B, D*
Official Answer: BD

janw Highly Voted 5 years, 2 months ago
correct = bc
upvoted 14 times

asefa Highly Voted 5 years, 2 months ago
i think B C r correct
upvoted 8 times

gio Most Recent 2 years, 1 month ago
Selected Answer: BC
100% BC
upvoted 2 times

thetuk05 2 years, 1 month ago
Selected Answer: BD
B. It runs automatically after each AWR snapshot. D. Results are written to the alert log.
upvoted 1 times

raferen10 2 years, 5 months ago
Selected Answer: BC
Correct:B,C

A. False. The time period analyzed by ADDM is defined by the last two snapshots (the last hour by default). B. True. An ADDM analysis task is performed and its findings and recommendations stored in the database every time an AWR snapshot is taken provided the STATISTICS_LEVEL parameter is set to TYPICAL or ALL C. True. DBA can run manual ADDM , for example in order to increase de period analyzed (default is time beetween two last AWR snaps) D. False. stores in AWR not in alert logfile.Data is stored both in memory and in the database E. False. The time period analyzed by ADDM is defined by the last two snapshots (the last hour by default).
upvoted 2 times

jareach 2 years, 6 months ago
Selected Answer: BC
I believe so
upvoted 1 times

477267 2 years, 7 months ago
BC is correct
upvoted 1 times

JESUSBB 2 years, 10 months ago
B and C are correct
upvoted 1 times

Patrick9230 3 years, 1 month ago
Selected Answer: BC
BC is correct
upvoted 1 times

asefa 5 years, 2 months ago

ill go with BC

upvoted 4 times

Jatindra 5 years, 2 months ago

BC is correct. Results are written to the alert log. -Incorrect .. stores in AWR not in alert logfile.

upvoted 4 times

janw 5 years, 2 months ago

The time period analyzed by ADDM is defined by the last two snapshots (the last hour by default).

upvoted 1 times

janw 5 years, 2 months ago

An ADDM analysis task is performed and its findings and recommendations stored in the database every time an AWR snapshot is taken provided the STATISTICS_LEVEL parameter is set to TYPICAL or ALL

upvoted 3 times

Question #8 *Topic 1*

Which two are true about server-generated alerts? (Choose two.)

A. Stateful alerts must be created by a DBA after resolving the problem.

B. Stateless alerts can be purged manually from the alert history.

C. Stateless alerts can be cleared manually.

D. Stateless alerts are automatically cleared.

E. Stateful alerts are purged automatically from the alert history.

Suggested Answer: *A, C*

Reference: https://jameshuangsj.wordpress.com/2019/12/01/clears-stateless-alerts-in-oem-by-using-emcli/

Community vote distribution



janw Highly Voted 5 years, 2 months ago

typo. correct: BC

upvoted 11 times

xRodge Highly Voted 4 years, 1 month ago

Best to start with defining stateful/stateless: A stateful alert is a threshold based alert (ie, tablespace usage above 90%), a stateless alert is non-threshold based alert (ie, a capture aborted with an ORA- error). A stateful alert first appears in DBA_OUTSTANDING_ALERTS. When cleared, it goes to DBA_ALERT_HISTORY. A stateless alert goes straight to DBA_ALERT_HISTORY

A stateless alert can be considered as a "point in time" error, the DB tells you about it and then forgets that it ever spoke about it.

Therefore, each answer:

A: FALSE - the DB creates stateful alerts, not a DBA. B: TRUE - a DBA can perform a purge to bulk remove alerts, perhaps based on date C: TRUE - a DBA can clear individual alerts D: FALSE - Stateless alerts are never checked again, by the DB, so cannot be automatically cleared E: FALSE - Stateful alerts are purged from the "Outstanding Alerts" and put into the alert history, where they will remain.

Therefore, BC are correct

upvoted 10 times

auwia 1 year, 10 months ago

E is false but it's wrong your explanation. Clearing an alert send it to Alert History Purging an alert removes it from the Alert History.

upvoted 1 times

auwia 1 year, 10 months ago

Selected Answer: BE

A. Stateful alerts must be created by a DBA after resolving the problem. --> False, they are created automatically. B. Stateless alerts can be purged manually from the alert history. --> True C. Stateless alerts can be cleared manually. --> False, *MUST be cleared manually. D. Stateless alerts are automatically cleared. --> False, you can only clear manually. E. Stateful alerts are purged automatically from the alert history. --> True.

upvoted 3 times

auwia Most Recent 1 year, 10 months ago

Proof: https://docs.oracle.com/en/database/oracle/oracle-database/19/admin/monitoring-the-database.html#GUID-3B999A58-21A7-40FB-A36E-2A113A83F2CF

https://docs.oracle.com/cd/E73210_01/EMADM/GUID-D0428ED9-FCA8-44D6-AB5B-B8FF4139BF0D.htm#EMADM12104

upvoted 1 times

auwia 1 year, 10 months ago

Selected Answer: BC

B and C are corrects.

upvoted 1 times

auwia 1 year, 10 months ago

I was wrong, C is false because it is a MUST and not that you can clear manually. You can only clear manually, so the verb has to be MUST and not CAN.

upvoted 1 times

kaka321 2 years, 1 month ago

BC is write see https://docs.oracle.com/cd/E73210_01/EMADM/GUID-D0428ED9-FCA8-44D6-AB5B-B8FF4139BF0D.htm#EMADM12104 section 2.4.3 Clearing Stateless Alerts for Metric Alert Event Types For metric alert event types, an event (metric alert) is raised based on the metric threshold values. These metric alert events are called stateful alerts. For those metric alert events that are not tied to the state of a monitored system (for example, snapshot too old, or resumable session suspended), these alerts are called stateless alerts. Because stateless alerts are not cleared automatically, they need to be cleared manually. You can perform a bulk purge of stateless alerts using the clear_stateless_alerts EM CLI verb.

upvoted 1 times

gio 2 years, 4 months ago

Selected Answer: CE

"TO PURGE" IS NOT "TO CLEAR". "Background processes periodically flush the data to the Automatic Workload Repository to capture a history of metric values. The alert history table and ALERT_QUE are purged automatically by the system at regular intervals." https://docs.oracle.com/en/database/oracle/oracle-database/19/admin/monitoring-the-database.html#GUID-3B999A58-21A7-40FB-A36E-2A113A83F2CF

upvoted 1 times

yumyummyyum 2 years, 4 months ago

I wonder whether the answer should be BD. As stateless alerts go directly to DBA_ALERT_HISTORY, how can C be correct as you wont be able to find that out for clearing manually. Pls help to correct me if I'm wrong.

upvoted 1 times

raferen10 2 years, 5 months ago

Selected Answer: BC

Correct: B,C

A. False. A stateful alert is a threshold based alert (ie, tablespace usage above 90%),the DB creates stateful alerts, not a DBA. B. True. DBA can perform a purge to bulk remove alerts, perhaps based on date C. True. DBA can clear individual alerts D. False. Stateless alerts are never checked again, by the DB, so cannot be automatically cleared. Only resolved alerts are cleared automatiquely E. False. Stateful alerts are purged from the "Outstanding Alerts"(DBA_OUTSTANDING_ALERTS) and put into the alert history, where they will remain.

upvoted 1 times

vkra 2 years, 6 months ago

Selected Answer: BC

see comment from xRodge

upvoted 1 times

jareach 2 years, 6 months ago

Selected Answer: CE

Stateful alerts are purged after 24 hours

upvoted 1 times

JESUSBB 2 years, 10 months ago

B and C are correct

upvoted 1 times

Patrick9230 3 years, 1 month ago

Selected Answer: BC

BC is correct

upvoted 1 times

Neil107 4 years ago

It's BC

upvoted 1 times

Jatindra 4 years, 5 months ago

Correct Ans: BC There are two kinds of server-generated alerts: threshold and nonthreshold. Most server-generated alerts are configured by setting a warning and critical threshold values on database metrics. You can define thresholds for more than 120 metrics, including the following: • Physical Reads Per Sec • User Commits Per Sec • SQL Service Response Time Except for the Tablespace Space Usage metric, which is database related, the other metrics are instance related. Threshold alerts are also referred to as stateful alerts, which are automatically cleared when an alert condition clears. Stateful alerts appear in DBA_OUTSTANDING_ALERTS and, when cleared, go to DBA_ALERT_HISTORY. Other server-generated alerts

correspond to specific database events, such as ORA-* errors, “Snapshot too old” errors, Recovery Area Low On Free Space, and Resumable Session Suspended. These are non-threshold-based alerts, also referred to as stateless alerts. Stateless alerts go directly to the History table

upvoted 2 times

RinD 4 years, 7 months ago

B and C

upvoted 2 times

ace03 4 years, 7 months ago

C and E

upvoted 1 times

Jatindra 5 years, 2 months ago

Correct Ans: B,C Student Guide: Except for the tablespace space usage metric, which is database related, the other metrics are instance related. Threshold alerts are also referred to as stateful alerts which are automatically cleared when an alert condition clears. Stateful alert appears in DBA_OUTSTANDING_ALERTS and when cleared go to DBA_ALERT_HISTORY. Other server-generated alerts correspond to specific database events such as ORA-* errors, "Snapshot too old" errors, Recovery Area Low on Free Space, Resumable Session Suspended. These are non threshold based alerts, also referred to as stateless alerts. Stateless alerts go directly to the History table. +++ Most alerts (such as "Out of Space") are cleared automatically when the cause of the problem disappears. However, other alerts (such as generic alert log errors) are sent to you for notification and must be acknowledged by you. After taking the corrective measures, you acknowledge an alert by clearing or purging it. Clearing an alert sends the alert to the Alert History which is accessible from Monitoring sub menu. Purging an alert removes it from the Alert History.

upvoted 3 times

Question #9 *Topic 1*

Which three are located by using environment variables? (Choose three.)

A. the Optimal Flexible Architecture (OFA) compliant path to store Oracle software and configuration files.

B. the location of Oracle Net Services configuration files

C. the list of a disk group names to be mounted by an Oracle Automatic Storage Management (ASM) instance at startup

D. default directories for temporary files used by temporary tablespaces

E. the temporary disk space used by Oracle Installer during installation

F. the maximum number of database files that can be opened by a database instance

Suggested Answer: *B, D, F*
Official Answer: BDF

Community vote distribution
ABE (100%)

janw Highly Voted 5 years, 2 months ago
correct= A, B, E
upvoted 14 times

yurijk Highly Voted 4 years, 11 months ago
A - ORACLE_BASE, ORACLE_HOME - correct B - TNS_ADMIN, old TWO_TASK, etc - correct C - spfile parameter ASM_DISKGROUPS, not env variable D - [s]pfile parameter DB_CREATE_FILE_DEST, not env variable E - TMP, TEMP - correct F - [s]pfile parameter DB_FILES, not env variable
upvoted 11 times

gio Most Recent 2 years ago
Selected Answer: ABE
100% ABE
upvoted 1 times

raferen10 2 years, 5 months ago
Selected Answer: ABE
Correct: A,B,E

A. True. $ORACLE_HOME and $ORACLE_BASE B. True. $TNS_ADMIN (old $TWO_TASK) C. False. spfile parameter ASM_DISKGROUPS D. False. [s]pfile parameter DB_CREATE_FILE_DEST, not env variable E. True. TMP, TEMP - directories tah oracle use to storage temporal installation files F. False. Is inizialization parameter DB_FILES
upvoted 1 times

Patrick9230 3 years, 1 month ago
Selected Answer: ABE
ABE is correct
upvoted 3 times

SMagny 4 years, 10 months ago
Correct answer : A,B,E. D and F are incorrect.
upvoted 3 times

IOracle 5 years, 1 month ago
Correct B - TNS_ADMIN, C - The ASM_DISKGROUPS initialization parameter specifies a list of disk group names that an Oracle ASM instance mounts at startup when the SQL ALTER DISKGROUP ALL MOUNT statement is issued. D- DB_CREATE_FILE_DEST -> DB_CREATE_FILE_DEST initialization parameter in your initialization parameter file to identify the default location for the database server to create:

Data files

Temp files

Redo log files

Control files

Block change tracking files

upvoted 1 times

taotsumiau 5 years, 1 month ago

I agree with janw ABE

upvoted 1 times

nyanyanyao0826 5 years, 1 month ago

corrent = A, B, D

upvoted 1 times

logica 5 years, 1 month ago

Correct: ABD

upvoted 1 times

logica 5 years, 1 month ago

I agree with janw ABE

upvoted 1 times

janw 5 years, 2 months ago

A = $ORACLE_BASE

upvoted 1 times

janw 5 years, 2 months ago

D = control file

upvoted 1 times

janw 5 years, 2 months ago

CF = spfile

upvoted 1 times

janw 5 years, 2 months ago

E = TMP env variable

upvoted 1 times

janw 5 years, 2 months ago

B = TNS_ADMIN env variable

upvoted 1 times

Question #10 *Topic 1*

Which three are true about opatchauto? (Choose three.)

A. It performs a shutdown and then a restart of all processes in both Oracle Grid Infrastructure and Oracle Database home during the patching process.

B. It must be invoked by a user with root user privileges.

C. Patches are applied via opatchauto.

D. Users must always input patch plans to opatchauto.

E. It requires the Oracle Grid Infrastructure and Oracle Database instances to be shut down before being invoked.

F. It applies patches in nonrolling mode by default.

G. It is used to apply interim patches to Oracle Grid Infrastructure and Oracle Database home combinations.

Suggested Answer: *B, C, E*
Official Answer: BCE

RinD Highly Voted 4 years, 7 months ago

A. It performs a shutdown and then a restart of all processes in both Oracle Grid Infrastructure and Oracle Database home during the patching process. => True B. It must be invoked by a user with root user privileges. => true C. Patches are applied via opatchauto. => not always. D. Users must always input patch plans to opatchauto. => false, patch plans is optional E. It requires the Oracle Grid Infrastructure and Oracle Database instances to be shut down before being invoked. => false F. It applies patches in nonrolling mode by default. => false, default is rolling mode. G. It is used to apply interim patches to Oracle Grid Infrastructure and Oracle Database home combinations. => true

Correct: A,B,G Reference: https://docs.oracle.com/cd/E24628_01/doc.121/e39376/opatchauto_commands.htm#OPTCH585
https://docs.oracle.com/cd/E24628_01/doc.121/e39376/configuration_patching.htm#OPTCH149

upvoted 17 times

gio 2 years ago

Selected Answer: ABG

I agree with RinD

upvoted 1 times

casfdsaf 3 years, 4 months ago

Selected Answer: ABG

ADC correct

upvoted 2 times

casfdsaf 3 years, 4 months ago

ABG is right

upvoted 2 times

Steve26NL 3 years ago

B yes sure you run opatchauto with root user have you ever worked with Oracle before?

upvoted 1 times

Neil107 4 years ago

ABG, I think.

upvoted 2 times

Abdou1968 4 years, 10 months ago

A C & D are correct

upvoted 1 times

yurijk 4 years, 11 months ago

A B C (applied via opatch) G

upvoted 1 times

auwia Most Recent 1 year, 10 months ago

so C is wrong! ;)

upvoted 1 times

lOracle 5 years, 1 month ago

OPatchAuto must be run from the GI Home as a root user OPatchAuto apply command applies the patch Rolling Mode (Default Mode): When performing patching in Rolling mode, the ORACLE_HOME processes on a particular node are shut down, the patch is applied, then the node is brought back up again. This process is repeated for each node in the GI or RAC environment until all nodes are patched. Opatchauto is used also to apply interim patches A,B,C,G

upvoted 1 times

julica 5 years, 1 month ago

correct: ABG

upvoted 2 times

Jatindra 5 years, 2 months ago

Student Guide: opatchauto: Is a utility for applying interim patches to Oracle Grid Infrastructure and Database Home combinations. Usually found in the directory $ORACLE_HOME/OPatch Updated with Patch 6880880 for all product along with opatch Invoke with the command opatchauto after adding $ORACLE_HOME/OPatch to the PATH Invoked as the root user Will shutdown, patch and restart all processes in both the Oracle Grid Infrastructure and Database Home. Allow a single command to patch a Real Application Cluster (RAC) system or a Single Instance High Availability (SIHA) configuration.

upvoted 4 times

Jatindra 5 years, 2 months ago

A. It performs a shutdown and then a restart of all processes in both Oracle Grid Infrastructure and Oracle Database home during the patching process. ==> correct B. It must be invoked by a user with root user privileges. ==> it must be invoked by root.. C. Patches are applied via opatchauto. This is Correct but E is more accurate D. Users must always input patch plans to opatchauto. E. It requires the Oracle Grid Infrastructure and Oracle Database instances to be shut down before being invoked. incorrect F. It applies patches in nonrolling mode by default. G. It is used to apply interim patches to Oracle Grid Infrastructure and Oracle Database home combinations. Correct Correct: A, B, C, G (seems four correct options)

upvoted 1 times

amorimleandro Highly Voted 5 years ago

C is incorrect, the patches are applied via opatch. Opatchauto is just a tool: " A patch orchestration tool that generates patching instructions specific to your target configuration and then uses OPatch to perform the patching operations without user intervention. Specifically, OPatchAuto can: 1. Perform pre-patch checks. 2. Apply the patch 3. Start and stop the running servers. 4. Perform post-patch checks. 5. Roll back patches when patch deinstallation is required."

upvoted 5 times

Question #11 *Topic 1*

Which two are true about the character sets used in an Oracle database? (Choose two.)

A. Single-byte character sets provide better performance than multibyte character sets.

B. Unicode enables information from any language to be stored using a single character set.

C. Unicode is the only supported character set for Oracle databases created using Database Configuration Assistant (DBCA).

D. Single-byte character sets always use 7-bit encoding schemes.

E. Multibyte character sets allow more efficient space utilization than single byte character sets.

F. Single-byte character sets always use 8-bit encoding schemes.

Suggested Answer: *A, B*

Reference: https://docs.oracle.com/database/121/NLSPG/ch2charset.htm#NLSPG166

Community vote distribution

AB (100%)

ama Highly Voted 4 years, 11 months ago

Well this is clear that A, B are correct

A > Single-byte character sets result in better performance than multibyte character sets, and they also are the most efficient in terms of space requirements. Ref: https://docs.oracle.com/en/database/oracle/oracle-database/19/nlspg/choosing-character-set.html#GUID-BF26E01D-AB92-48FC-855A-69A5B3AF9A92

B> Unicode is a universal encoded character set that enables information from any language to be stored using a single character set.

Ref:https://docs.oracle.com/cd/B19306_01/server.102/b14225/ch6unicode.htm#:~:text=Unicode%20is%20a%20universal%20encoded%20character%20set%20that,been%20adopted%20by%20many%20software%20and%20hardware%20vendors.

upvoted 14 times

yurijk Highly Voted 4 years, 11 months ago

A "Single-byte character sets result in better performance than multibyte character set" (Logic & Oracle 12.2 Docs Application Development)

B "Unicode is a universal encoded character set that enables information from any language to be stored using a single character set" (Oracle 10.2 docs)

upvoted 5 times

gio Most Recent 2 years ago

Selected Answer: AB

100% AB

upvoted 1 times

raferen10 2 years, 5 months ago

Selected Answer: AB

Correct: A, B

A. True. Single-byte character sets result in better performance than multibyte character sets, and they also are the most efficient in terms of space requirements. B. True. Unicode is a universal encoded character set that enables information from any language to be stored using a single character set. C. False. you can chose other characters sets D. False. Can use 7-bits and 8-bits encoding schemas E. False. Single-byte character sets result in better performance than multibyte character sets, and they also are the most efficient in terms of space requirements. F. False.Can use 7-bits and 8-bits encoding schemas

upvoted 1 times

hilaire 2 years, 8 months ago

AB correct

upvoted 1 times

JESUSBB 2 years, 10 months ago

A and B correct

upvoted 1 times

Qwerus 2 years, 11 months ago

Selected Answer: AB

AB - correct

upvoted 2 times

Neil107 4 years ago

It's AB.

upvoted 2 times

asefa 4 years, 10 months ago

according to student guide, A and B are correct

upvoted 2 times

Abdou1968 4 years, 10 months ago

AB Correct

upvoted 2 times

cerebro2000x 5 years ago

AE single byte characters are space inefficient and multibyte characters are fast [Japanese are checked whether they aer 1,2,3,4 bytes in utf8 but are stored internally as uniform of 2 bytes in utf16

upvoted 1 times

git17 5 years, 2 months ago

A is correct then what is the another option is correct JANW?

upvoted 1 times

veginha 5 years, 2 months ago

i think B is the other one, Unicode is intended to include every character that will ever be required.

upvoted 3 times

mihanik 5 years ago

Voted for AB https://docs.oracle.com/database/121/NLSPG/ch2charset.htm#NLSPG175

upvoted 1 times

janw 5 years, 2 months ago

B: incorrect You should generally select the Unicode character set AL32UTF8, because it supports most languages of the world.

C: incorrect Starting from Oracle Database 12c Release 2, if you use Oracle Universal Installer or Database Configuration Assistant (DBCA) to create the database, then the default database character set is AL32UTF8.

A: correct Single-byte character sets result in better performance than multibyte character sets, and they also are the most efficient in terms of space requirements.

E: incorrect Single-byte character sets result in better performance than multibyte character sets, and they also are the most efficient in terms of space requirements.

upvoted 1 times

janw 5 years, 2 months ago

Single-byte encoding schemes are classified as one of the following types:

7-bit encoding schemes 8-bit encoding schemes

DF: incorrect

upvoted 1 times

Question #12 *Topic 1*

Which three are true about monitoring waits for sessions and services? (Choose three.)

A. V$SESSION_EVENT displays all waits for all past and existing sessions if the wait has occurred at least once for a session.

B. V$SERVICE_EVENT displays all waits for all services if the wait has occurred at least once for a service.

C. V$SESSION_WAIT_CLASS displays waits broken down by wait class only for waiting sessions.

D. V$SESSION_WAIT and V$SESSION both contain details of the event on which a non-waiting session last waited.

E. V$SESSION_EVENT displays all waits for all past sessions if the wait has occurred at least once for a session.

F. V$SESSION_WAIT and V$SESSION both contain details of the event on which a session is currently waiting.

Suggested Answer: *B, D, F*

Official Answer: BDF

Community vote distribution



brolem Highly Voted 4 years, 3 months ago

Answer: B, D, F Explanation : A, E - False : V$SESSION_EVENT displays information on waits for an event by a session. V$SESSION_WAIT_HISTORY displays the last 10 wait events for each active session. B - V$SERVICE_EVENT displays aggregated wait counts and wait times for each wait statistic. C - False : V$SESSION_WAIT_CLASS displays the time spent in various wait event operations on a per-session basis. D, F - True : V$SESSION_WAIT displays the current or last wait for each session. V$SESSION.EVENT If the session is currently waiting, then the resource or event for which the session is waiting. If the session is not in a wait, then the resource or event for which the session most recently waited.

Resource : https://docs.oracle.com/en/database/oracle/oracle-database/19/refrn/V-SESSION_EVENT.html#GUID-5723D404-83F7-4C8D-8BF2-031B3FFEAAA8

upvoted 14 times

xRodge Highly Voted 4 years, 7 months ago

The session/service views show what is happening NOW. You can append _HISTORY onto these (eg, V_SESSION_WAIT_HISTORY) to see what happened in the past, to sessions that have disconnected (on disconnect, Oracle will, basically, copy that session's v$ data into the corresponding v$_HISTORY views. Therefore, any answer that refers to anything in the past, here, must be wrong. Another note is that all sessions are ALWAYS waiting for something. An "active" wait would be waiting for I/O to return. An "idle" wait would be waiting for the user to ask the DB to do something. They're both considered to be "waits", so the DB is always waiting.

A: FALSE - It shows these details for existing sessions, not past sessions. B: TRUE - The definition is exactly what the view does C: TRUE - As mentioned, all sessions are always waiting. So, this is true D: FALSE - There is no such thing as a "non-waiting session", so this must be false. E: FALSE - This references older sessions, so you would need to access the _HISTORY view for this data. F: TRUE - This is exactly what these views do.

Therefore, BCF are correct

upvoted 9 times

Alvinzzz 2 years, 8 months ago

C is incorrect as the view is grouping the similar waits into high-level wait class for the session.

upvoted 1 times

Datajimm 3 years, 10 months ago

C is false, since it "displays the time spent in various wait event operations on a per-session basis", therefore it shows no matter if the session is in wait or not https://docs.oracle.com/en/database/oracle/oracle-database/19/refrn/V-SESSION_WAIT_CLASS.html#GUID-B890F027-CE8A-452B-96B2-213985B21484

D is true: "If the session is currently waiting, then the resource or event for which the session is waiting. If the session is not in a wait, then the resource or event for which the session most recently waited." https://docs.oracle.com/en/database/oracle/oracle-database/19/refrn/V-SESSION_WAIT.html#GUID-4EDAB293-F3FC-40FE-BC75-4FEE6A4D7705

upvoted 4 times

Haeusi Most Recent 11 months, 3 weeks ago

Selected Answer: BDF

C is wrong because of "only for waiting sessions". V$SESSION_WAIT_CLASS groups wait events by wait class (e.g., User I/O, System I/O, Concurrency) for current sessions. It shows cumulative wait statistics for each wait class, not just for sessions currently waiting.

upvoted 1 times

dancymonkey 2 years ago

BDF

D is incorrect, I think not only displays waits broken down by wait class V$SESSION_WAIT_CLASS displays the time spent in various wait event operations on a per-session basis. https://docs.oracle.com/en/database/oracle/oracle-database/19/refrn/V-SESSION_WAIT_CLASS.html#GUID-B890F027-CE8A-452B-96B2-213985B21484

upvoted 1 times

dancymonkey 2 years ago

sorry I meant C is incorrect

upvoted 1 times

LEOC71 2 years, 3 months ago

Selected Answer: BCF

The V$SERVICE_EVENT view shows, by service, the total waits for a particular event since instance startup. => B: TRUE The V$SESSION_WAIT_CLASS view aggregates the waits by session and wait class, since instance startup. All sessions are always waiting. => C: TRUE The V$SESSION_WAIT view lists the resources or events for which active sessions are waiting. The V$SESSION view also includes the current wait information. => F: TRUE & D: FALSE

upvoted 1 times

Guhborges 2 years, 7 months ago

Selected Answer: BDF

Sorry for the duplicate comment, but the right question is B,D,F

upvoted 1 times

Guhborges 2 years, 7 months ago

Selected Answer: BDF

B,D,F 100%

upvoted 2 times

Guhborges 2 years, 7 months ago

Selected Answer: BCF

B,C,F for sure.

upvoted 1 times

gio 2 years, 10 months ago

Selected Answer: BDF

both V$SESSION_WAIT and V$SESSION contains information of actual wait or last wait (example: column "state" shows "waiting" or "waited for short"). C incorrect because contains statistical counts for non-waiting sessions too.

upvoted 1 times

gio 2 years, 6 months ago

I'm not sure because they doesn't show "details". I think ABE can be right.. why not??

upvoted 1 times

vkra 3 years ago

Selected Answer: BDF

Answer: B D F

upvoted 1 times

Patrick9230 3 years, 7 months ago

Selected Answer: BCF

BCF is correct

upvoted 2 times

vska 3 years, 10 months ago

BDF Sessions are not always waiting - they can be in working state (ON CPU). Check description of EVENT column: If the session is currently waiting, then the resource or event for which the session is waiting. If the session is not in a wait, then the resource or event for which the session most recently waited.

upvoted 4 times

Neil107 4 years, 6 months ago

It's BCF. You're checking for current waits, not past waits.

upvoted 2 times

nobody347 4 years, 6 months ago

I think C is wrong The V$SESSION_WAIT_CLASS view allows you to see the session wait information broken down by wait class for each session.
https://oracle-base.com/articles/10g/performance-tuning-enhancements-10g

upvoted 5 times

Abdou1968 5 years, 4 months ago

form me BCF are corrects

upvoted 3 times

janw 5 years, 6 months ago

E and A incorrect? as it speaks about past sessions, i presume AWR is necessary (e.g. hist tables)

upvoted 1 times

Alejandrrro 5 years, 1 month ago

Yes, I came to the same conclusion about A and E being incorrect. C also looks incorrect to me as v$session_wait_class will also show the waits happened for non waiting sessions, for example 'Idle'. So I got B,D,F as correct answers for now. What do you think?

upvoted 8 times

janw 5 years, 6 months ago

C incorrect: v$session_wait_class will also show the waits happened for non waiting sessions

upvoted 2 times

Question #13 *Topic 1*

You must transport the UNIVERSITY tablespace from one database to another.
The UNIVERSITY tablespace is currently open read/write.
The source and destination platforms have different endian formats.
Examine this list of actions:
1. Make the UNIVERSITY tablespace read-only on the source system.
2. Export the UNIVERSITY tablespace metadata using EXPDP.
3. Convert the UNIVERSITY tablespace data files to the destination platform format using RMAN on the source system.
4. Copy the UNIVERSITY tablespace data files to the destination system.
5. Copy the Data Pump dump set to the destination system.
6. Convert the UNIVERSITY tablespace data files to the destination platform format using RMAN on the destination system.
7. Import the UNIVERSITY tablespace metadata using IMPDP.
8. Make the UNIVERSITY tablespace read/write on the destination system.
Which is the minimum number of actions required, in the correct order, to transport the UNIVERSITY tablespace?

A. 1, 2, 4, 5, 7, 8

B. 1, 2, 4, 6, 7, 8

C. 1, 2, 3, 4, 5, 7, 8

D. 1, 2, 3, 4, 5, 6, 7, 8

E. 2, 4, 5, 6, 7

Suggested Answer: *B*
Official Answer: B

Community vote distribution



asefa Highly Voted 5 years, 4 months ago
easy way we must convert either on destination or source so A is incorrect (no option 3 and 6) we must copy data pump set from source to destination so B is incorrect (no step 5) no need to convert on destination if we aleady did conversion on source so D is incorrect (both step 3 and 6) are included we must make tablespace read only son E is incorrect so answer must be C
upvoted 15 times

egore0496 4 years ago
Why not B According https://docs.oracle.com/en/database/oracle/oracle-database/19/admin/transporting-data.html#GUID-E4C56852-73A5-44A2-BB10-938831DA6E4C C but what about network mode of data pump or nfs/cifs export?
upvoted 2 times

egore0496 4 years ago
15.2.4 Transporting a Database Over the Network https://docs.oracle.com/en/database/oracle/oracle-database/19/admin/transporting-data.html#GUID-032C9715-B2F8-4ACD-8A1C-C1A899DEA3C1
upvoted 2 times

asduta 3 years, 8 months ago
thats not in the choices isnt it?
upvoted 1 times

janw Highly Voted 5 years, 6 months ago
agree with c must contain 1, 5, 8 must contain 3 and 6 (not both)
upvoted 9 times

Alejandrrro 5 years, 1 month ago
Agree with C too MOS Doc ID 371556.1 describes the action plan almost the same as in C. Files are converted on the source system
upvoted 2 times

alrech Most Recent 1 year, 7 months ago
Selected Answer: C

The correct order is 3, 1, 2, 5, 4, 6, 7, 8. Reference: https://mylearn.oracle.com/ou/ekit/86212/38560/dc9de5cb-8c91-4bd5-bc30-18406127a61b/course Page: 296

upvoted 1 times

jaimegarcia 2 years, 2 months ago

Selected Answer: B

yes , transportable tablespace

upvoted 1 times

LEOC71 2 years, 3 months ago

Selected Answer: C

Step 2 is part of all alternatives, then the data pump set created must be transferred to the destination system (step 5). => C: TRUE B:FALSE => to be true step 2 shouldn't be mentioned as part of the procedure, only step 7 (using impdp network mode).

upvoted 1 times

flaviogcmelo 2 years, 5 months ago

Selected Answer: C

In my opinion, choosing alternative C instead of B is the detail of making the dump file available to the destination database. https://docs.oracle.com/en/database/oracle/oracle-database/19/spmds/migrating-oracle-database.html#GUID-F224185E-6AEC-45FC-B4F3-7F8B6469606D

upvoted 1 times

gio 2 years, 6 months ago

Selected Answer: C

I think C

upvoted 1 times

Guhborges 2 years, 7 months ago

Selected Answer: C

For me, the right answer should B, but here, we have a lot of person saying the correct one is C, so go with the others.

upvoted 1 times

Marcello86CT 3 years, 4 months ago

I would say that correct Answer is : E

With Data Pump Conventional Export and Import we no need to make any convesrion because it support different endian format. In conventional mode is not requeired neither to set db in read only.

upvoted 1 times

abdelouahab 3 years, 3 months ago

Hi,

the conversion is a must between different endianess see https://docs.oracle.com/en/database/oracle/oracle-database/19/spmdu/task-4-transport-the-tablespace-set.html#GUID-41D10DCA-7733-4E4C-A41A-B94F63A82F7C

upvoted 1 times

Patrick9230 3 years, 7 months ago

Selected Answer: C

Must C

upvoted 3 times

egore0496 4 years ago

B 5 is not must https://docs.oracle.com/en/database/oracle/oracle-database/19/sutil/oracle-data-pump-export-utility.html#GUID-72973E3D-FC0A-47E8-A62A-6DF8AD83138D

Network Considerations for Oracle Data Pump Export

upvoted 2 times

asduta 3 years, 8 months ago

5 is a must since above choices doesn't list other type of moving the dump file, including one you mentioned.

upvoted 1 times

Neil107 4 years, 6 months ago

It's C. It is always preferred/recommended to convert at the source, as far as possible, rather than convert it at the destination. Saves a lot of headache!

upvoted 1 times

Dhanushbh 4 years, 7 months ago

C is correct [Doc ID 2013271.1]

upvoted 1 times

fukaka 4 years, 8 months ago

B is correct

E is wrong: 1 is a must D is wrong: 3 & 6 need either one A & C are wrong: 5 is not a must if using network mode of data pump

upvoted 4 times

Jatindra 4 years, 11 months ago

Correct Answer is E

upvoted 1 times

Bda 4 years, 11 months ago

I also think that it should be B.

upvoted 1 times

Abdou1968 5 years, 4 months ago

B is a correct answer

upvoted 1 times

Question #14 *Topic 1*

Which two are true about OS groups and users for Oracle Grid Infrastructure and the Oracle Relational Database Management System (RDBMS)? (Choose two.)

A. By default, members of the OSASM group can access Automatic Storage Management and RDBMS instances.

B. The primary group for the Oracle Grid Infrastructure and Oracle Database owners must be the Oracle Inventory group.

C. The Oracle Grid Infrastructure installation must be owned by the grid user.

D. The Oracle Grid Infrastructure owner owns Oracle Restart and Oracle Automatic Storage Management binaries.

E. The Oracle Grid Infrastructure owner must have OSOPER, OSBACKUPDBA, and OSKMDBA as secondary groups.

F. The same OSDBA group must be used for Automatic Storage Management and the Oracle Database.

Suggested Answer: *E, F*
Reference: https://docs.oracle.com/database/121/CWWIN/usrgrps.htm#CWWIN-GUID-6FD3C6AE-5A99-4C6D-85D6-868CBA6F5DAE

ald85 Highly Voted 5 years, 2 months ago
I think B D
upvoted 14 times

Alejandrrro 4 years, 7 months ago
It's confirmed by the below quotes

https://docs.oracle.com/database/121/CWLIN/usrgrps.htm#CWLIN483

The -g option specifies the primary group, which must be the Oracle Inventory group. For example: oinstall.

You can designate a unique group for the SYSASM system privilege, separate from database administrator groups, or you can designate one group as the OSASM and OSDBA group, so that members of that group are granted the SYSASM and SYSDBA privilege to grant system privileges to administer both the Oracle ASM instances and Oracle Database instances.

If you are creating this user to own both Oracle Grid Infrastructure and an Oracle Database installation, then this user must have the OSDBA for ASM group as a secondary group.

In Oracle documentation, a user created to own the Oracle Grid Infrastructure binaries is called the grid user. This user owns both the Oracle Clusterware and Oracle Automatic Storage Management binaries.
upvoted 2 times

lOracle Highly Voted 5 years, 1 month ago
A - false - Members of this group are granted the SYSASM system privilege to administer Oracle ASM.SYSASM system privileges do not grant access privileges on an Oracle Database instance. B- Correct - C - false Only if you wish to define role separation but is not a must D - correct E- false OSOPER is optional is not a must F - the same OSDBA could be used but is not a must (usually SYSDBA and SYSASM are used for separation us duties)
upvoted 7 times

gio Most Recent 2 years ago
Selected Answer: BD
BD sure
upvoted 1 times

ScottL 2 years, 3 months ago
Selected Answer: BD
B D correct
upvoted 2 times

hilaire 2 years, 8 months ago
BD correct
upvoted 1 times

flaviogcmelo 2 years, 11 months ago

Selected Answer: BD

Based on 19c documentation. https://docs.oracle.com/en/database/oracle/oracle-database/19/cwlin/creating-users-groups-paths-for-grid-infrastructure.html#GUID-3D41DE23-ED6B-4AF1-BC7F-4EFA1D25E86C

upvoted 2 times

Patrick9230 3 years, 1 month ago

Selected Answer: BD

B D is correct

upvoted 1 times

arpz 3 years, 11 months ago

Agree with B and D is the answer

upvoted 1 times

ObserverPL 4 years, 3 months ago

by eliminating wrong answers - B and D are ok

upvoted 1 times

Abdou1968 4 years, 10 months ago

B D are correct

upvoted 1 times

Sha7 5 years, 1 month ago

Answer is B D

upvoted 1 times

Question #15 *Topic 1*

Which four are true about duplicating a database using Recovery Manager (RMAN)? (Choose four.)

A. Duplication can be done by having the auxiliary database instance pull backup sets from the target database instance.

B. A connection to an auxiliary instance is always required.

C. A subset of the target database can be duplicated.

D. A new DBID is always created for the duplicated database.

E. A connection to the recovery catalog instance is always required.

F. A backup of the target database is always required.

G. Duplication can be done by having the target database instance push copies to the auxiliary database instance.

H. A connection to the target database instance is always required.

Suggested Answer: *A, B, F, G*

Official Answer: ABFG

Community vote distribution

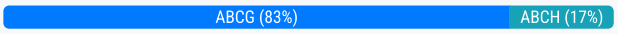

SMagny Highly Voted 4 years, 9 months ago

Answer : A,B,C,G

upvoted 16 times

janw Highly Voted 5 years, 2 months ago

correct = ABCG

upvoted 8 times

auwia Most Recent 1 year, 10 months ago

Selected Answer: ABCG

H is false, you can have: Backup-Based Duplication without a Target Database: https://docs.oracle.com/en/database/oracle/oracle-database/23/bradv/rman-duplicating-databases.html#GUID-D38B96D6-72C5-4760-8253-7D5C474C7914

upvoted 2 times

gio 2 years ago

Selected Answer: ABCG

DBID doesn't change if you move file without rman

upvoted 1 times

ScottL 2 years, 3 months ago

Selected Answer: ABCH

A,B,C,G correct

upvoted 1 times

JESUSBB 2 years, 10 months ago

A, B,C and G are correct

upvoted 1 times

Patrick9230 3 years, 1 month ago

Selected Answer: ABCG

A B C G is correct

upvoted 2 times

chefdba 3 years, 9 months ago

ABCG were correct answers in the exam

upvoted 1 times

Dhanushbh 4 years, 1 month ago

No need backup for active database duplicate so F incorrect . ABCG correct answers

upvoted 1 times

Abdou1968 4 years, 10 months ago

i think that ABCG

upvoted 5 times

asefa 5 years, 2 months ago

my answer ABCG

upvoted 6 times

Jatindra 5 years, 2 months ago

Student Guide: A duplicate database is a copy of your target database. With the FOR STANDBY clause, it keeps the same unique database identifier(DBID); If FOR STANDBY not specified it creates a new DBID.

The duplicate database can include the same content or only a subset from the source database. It can be in the same host or a separate host. The principal work of the duplication is performed by the auxiliary channels.These channels correspond to a server session on the auxiliary instance on the destination host for backup based duplication. For active database duplication the target channels perform the work of pushing data file copies to the auxiliary instance (if number of allocated target channels is greater than the number of allocated auxiliary channels).

upvoted 4 times

janw 5 years, 2 months ago

D: incorrect (duplicate standby keeps dbid)

upvoted 1 times

Abdou1968 4 years, 10 months ago

Wrong, in general, when using Rman duplicate, a new dbid is created

upvoted 1 times

Abdou1968 4 years, 10 months ago

so i thinh that ABCD seem also to be a a good answer

upvoted 1 times

navingupta52 4 years, 2 months ago

It says "always" in option D. It retains DBID if full backup is restored. If subset is restored new DBID gets generated. So D is incorrect.

upvoted 1 times

gio 2 years, 4 months ago

D is wrong because when you use "connection only to auxiliary" mode you can move manually needed files. When you manually move the files RMAN duplicate doesn't change DBID. So ABCG

upvoted 1 times

Question #16 *Topic 1*

A container database called CDB1 is OMF-enabled.
PDB_FILE_NAME_CONVERT is not configured in CDB1.
PDB1 was unplugged from CDB1 earlier in the week.
Examine this command, which will be executed in CDB1:
CREATE PLUGGABLE DATABASE pdb1 -
USING `~/u01/app/oracle/oradata/pdb1.xml'
SOURCE_FILE_NAME_CONVERT =
(`~/u01/app/oracle/oradata/', `~/u02/app/oracle/oradata/');
Which two are true? (Choose two.)

A. PDB1 data files already exist in the correct location.

B. DBMS_PDB.CHECK_PLUG_COMPATIBILITY must be run in CDB1 before executing the command.

C. PDB_FILE_NAME_CONVERT must be set before executing the command.

D. /u01/app/oracle/oradata/pdb1.xml does not contain the current locations of data files for PDB1.

E. PDB1 must be dropped from CDB1.

Suggested Answer: *C, D*
Reference: http://anjaniappsdba.blogspot.com/2016/09/oracle-database-12102c-hot-cloning-of.html

xRodge Highly Voted 4 years, 1 month ago
The question gives the impression that the source files were held at /u01/app/oracle/oradata, but are now physically located at /u01/app/oracle/oradata. The xml description file is still at /u01/app/oracle/oradata/pdb1.xml

So, for each: A: FALSE - SOURCE_FILE_NAME_CONVERT is only needed if the XML file does not accurately describe the current location of the pdb's files. B: FALSE - If pdb1 ran correctly before, it will run correctly now - assuming that no major upgrade of CDB1 has happened. C: FALSE - PDB_FILE_NAME_CONVERT is needed only when OMF is not used. The question tells us that OMF is enabled. D: TRUE - The use of the SOURCE_FILE_NAME_CONVERT parameter indicates that the DB files have been moved, so the xml file does not hold the correct locations E: TRUE - Unplugging a database will result in the PDB remaining in a MOUNT state. It is still part of the CDB and will still be backed up etc... After unplugging, a PDB needs to be dropped before a new PDB can be created with the same name.

Therefore, DE are correct
upvoted 11 times

Amaora Highly Voted 3 years, 11 months ago
The PDB must be dropped from the CDB before it can be plugged back into the same CDB. When the XML metadata file is specified, the XML file includes the full paths of the PDB files. When the .pdb archive file is specified, the XML metadata file contains the relative file names only.

https://docs.oracle.com/en/database/oracle/oracle-database/21/multi/plugging-in-a-pdb.html#GUID-3132C4E5-8734-458B-8B0C-EFC77A45FFA3

I believe A, E are correct
upvoted 5 times

Emperor11 Most Recent 1 year, 8 months ago
Hi people I just did the test, the correct answers are D and E, when you try to create a pdb with an xml it gives you error if you don't delete the pdb that was unplug so E is correct and the command is executed correctly when you specify the new destination of the data files with the SOURCE_FILE_NAME_CONVERT parameter if you don't do this it is not able to find any data file you can check it if you run a cat on the xml file you will see how the old path is there, good luck!
upvoted 3 times

auwia 1 year, 9 months ago
Selected Answer: DE
D and E are corrects. A is false because the data files "position" is specified in the .xml file and we don't know the content. So D is true, because for sure the SOURCE_FILE_NAME_CONVERT will change the current path from u01 to u02.

upvoted 1 times

gio 2 years, 4 months ago

Selected Answer: DE

when you use SOURCE_FILE_NAME_CONVERT you convert file path written in .xml file. So you change /u01/.. to /u02/.. So xml doesn't contain the current location (D is correct) PDB1 must be dropped because is in mount state after unplug (E is correct)

upvoted 2 times

jackymak 2 years, 9 months ago

Selected Answer: DE

I agree D,E

upvoted 1 times

Patrick9230 3 years, 1 month ago

Selected Answer: DE

A is mean the new create command output but exactly the new PDB1 have not create success

upvoted 1 times

Franky_T 3 years, 3 months ago

OK, so three possible answers here and not two.

A must be correct. The location of the data files used to be in '/u01/app/oracle/oradata/' when the PDB was unplugged, so if the SOURCE_FILE_NAME_CONVERT clause is used we are telling the CDB that the data files are now located in '/u02/app/oracle/oradata/', i.e. "already in the correct location". B is wrong. You would only use DBMS_PDB.CHECK_PLUG_COMPATIBILITY when you use a different CDB. C is wrong. You only need to use PDB_FILE_NAME_CONVERT when not using OMF. D is correct. Hence the fact that the SOURCE_FILE_NAME_CONVERT was included. E is correct. When a PDB is unplugged from a CDB, then the STATUS column in CDB_PDBS of the unplugged PDB will be UNPLUGGED. A PDB must be dropped from a CDB before it can be plugged back into the same CDB.

upvoted 2 times

freemun05 3 years, 4 months ago

WOW, not a single correct answer here))))))))))))

upvoted 1 times

Yema 3 years, 3 months ago

do you have the correct answer?

upvoted 1 times

casfdsaf 3 years, 4 months ago

Selected Answer: AE

AE correct

upvoted 2 times

chefdba 3 years, 9 months ago

A and E … 100% correct

upvoted 3 times

Aldrid 4 years, 2 months ago

AD is Correct. https://docs.oracle.com/en/database/oracle/oracle-database/19/multi/plugging-in-a-pdb.html#GUID-F23F5352-F456-48EE-AED7-71E3DEBAAED0

upvoted 1 times

Aldrid 4 years, 2 months ago

A. PDB1 data files already exist in the correct location. Correct: You use SOURCE_FILE_NAME_CONVERT to specify new correct location.

B. DBMS_PDB.CHECK_PLUG_COMPATIBILITY must be run in CDB1 before executing the command. WRONG: PDB1 was unplugged from CDB1 earlier. No NEED

C. PDB_FILE_NAME_CONVERT must be set before executing the command. WRONG: OMF is ENABLED.. Bust be if disabled

D. /u01/app/oracle/oradata/pdb1.xml does not contain the current locations of data files for PDB1.

CORRECT: We use SOURCE_FILE_NAME_CONVERT for correct path

E. PDB1 must be dropped from CDB1. WRONG. PDB1 already was unplugged

upvoted 1 times

Aldrid 4 years, 2 months ago

Switch to DE..

A: WRONG.. "Correct Location " is note accurate... We are using CONVERT to specify new correct...

E: Correct The PDB must be dropped from the CDB before it can be plugged back into the same CDB.

https://docs.oracle.com/en/database/oracle/oracle-database/18/multi/removing-a-pdb.html#GUID-3D9AA767-F03C-46EB-AF8E-6808AF1C9D9F

upvoted 2 times

navingupta52 4 years, 2 months ago

It will be AE. PDB1 files exist even when you unplug it and it doesn't get dropped so it remained in MOUNT state post unplugging.

See below:

SQL> show pdbs

CON_ID CON_NAME OPEN MODE RESTRICTED ---------- ------------------------------ ---------- ---------- 2 PDB$SEED READ ONLY NO 3 CRYPTO_PDB READ WRITE NO 4 PDB2 READ WRITE NO 5 APPCON1 READ WRITE NO

SQL> alter pluggable database pdb2 close;

SQL> ALTER PLUGGABLE DATABASE PDB2 UNPLUG INTO 'H:\oracle\Oracle19c\PDB2_info.xml' 2 ;

Pluggable database altered.

SQL> show pdbs

CON_ID CON_NAME OPEN MODE RESTRICTED ---------- ------------------------------ ---------- ---------- 2 PDB$SEED READ ONLY NO 3 CRYPTO_PDB READ WRITE NO 4 PDB2 MOUNTED 5 APPCON1 READ WRITE NO

Option A & E are correct.

upvoted 3 times

wiprooracle 4 years, 2 months ago

B,C looks correct

upvoted 2 times

RinD 4 years, 3 months ago

DE correct

upvoted 4 times

RinD 4 years, 3 months ago

datafile path: /u01/app/oracle/oradata/ is in xml file. But now datafile path need to be /u02/app/oracle/oradata/

upvoted 2 times

Shalom 4 years, 2 months ago

This SOURCE_FILE_NAME_CONVERT clause uses the files in the /disk2/oracle/pdb7 directory instead of the /disk1/oracle/pdb7 directory. In this case, the XML file describing a PDB specifies the /disk1/oracle/pdb7 directory, but the PDB should use the files in the /disk2/oracle/pdb7 directory.

SOURCE_FILE_NAME_CONVERT = ('/disk1/oracle/pdb7/', '/disk2/oracle/pdb7/')

upvoted 1 times

Rogazan 4 years, 3 months ago

Correct answer: AE

upvoted 2 times

Question #17 *Topic 1*

Which three are true about transporting databases across platforms using Recovery Manager (RMAN) image copies? (Choose three.)

A. By default, the transported database will use Oracle Managed Files (OMF)

B. Data files can be converted on the destination system.

C. Data files can be converted on the source system.

D. A new DBID is automatically created for the transported database.

E. Databases can be transported between systems with different endian formats.

F. The password file is automatically converted by RMAN.

Suggested Answer: *C, E, F*

Password file is automatically converted by RMAN.

Community vote distribution



lOracle Highly Voted 5 years, 7 months ago

A, B , C as per https://docs.oracle.com/en/database/oracle/oracle-database/18/multi/plugging-in-a-pdb.html#GUID-1691A7A2-372E-4E51-8A4F-15DAC6AB0FC3

upvoted 18 times

Heat418 3 years, 7 months ago

Yes, I think this is the correct answer. https://docs.oracle.com/en/database/oracle/oracle-database/19/bradv/rman-transporting-data-across-platforms.html#GUID-65AADCB6-CC9A-4229-9AB8-805C37E4471F ABC E is wrong: You can also use RMAN to transport an entire database to a different platform so long as the two platforms have the same endian format.

upvoted 3 times

GraceYu Highly Voted 5 years, 8 months ago

E seems incorrect. Transport database needs same endian format.

https://docs.oracle.com/en/database/oracle/oracle-database/18/bradv/rman-transporting-data-across-platforms.html#GUID-65AADCB6-CC9A-4229-9AB8-805C37E4471F

You can use RMAN to transport tablespaces across platforms with different endian formats. You can also use RMAN to transport an entire database to a different platform so long as the two platforms have the same endian format.

upvoted 7 times

Naku10 Most Recent 1 year, 7 months ago

Selected Answer: ABC

Note 2013540.1 (My Oracle Support): "When you transport entire database the source platform and the destination platform must use the same endian format."

upvoted 1 times

antonica 1 year, 10 months ago

B, is not possible, You can perform tablespace conversion with the RMAN CONVERT TABLESPACE command on the source host, but not on the destination host. https://docs.oracle.com/en/database/oracle/oracle-database/19/bradv/rman-transporting-data-across-platforms.html#GUID-5515B50B-2FCF-4C94-BEA3-16070E4E2334

upvoted 1 times

dancymonkey 2 years ago

BCE

BC When you use the RMAN CONVERT command to convert data, you can either convert the data on the source platform after running Data Pump export, or you can convert the data on the target platform before running Data Pump import. In either case, you must transfer the data files from the source system to the target system.

E - Converting Data Files on the Target System Before Import... The endianness of the source platform is different from the endianness of the target platform.

https://docs.oracle.com/en/database/oracle/oracle-database/19/spmds/converting-data-using-rman.html#GUID-DDFA71A3-396A-440A-B9AB-E0970F843904:~:text=When%20you%20use,the%20target%20system.

upvoted 1 times

guretto 2 years, 2 months ago

https://docs.oracle.com/en/database/oracle/oracle-database/18/bradv/rman-transporting-data-across-platforms.html#GUID-D4C86E36-8EF8-489D-84CE-B8F045C8AB20 Par 28.2.2 "... 28.2.2 Overview of Database Conversion Using Image Copies To convert a whole database to a different platform, both platforms must use the same endian format..."

upvoted 1 times

LEOC71 2 years, 2 months ago

Selected Answer: ABC

E: False. Note 2013540.1 (My Oracle Support): "When you transport entire database the source platform and the destination platform must use the same endian format."

upvoted 1 times

zouve 2 years, 5 months ago

Selected Answer: BCE

A is wrong. depends on the specific configuration of the Oracle Database environment, and it's not directly related to the process of transporting a database using RMAN.

E is true. with the appropriate commands RMAN> CONVERT DATAFILE '/path/tbs_31.f', '/path/tbs_32.f', '/path/tbs_41.f' TO PLATFORM="Solaris[tm] OE (32-bit)" FROM PLATFORM="HP TRu64 UNIX" DB_FILE_NAME_CONVERT= "/path_source/", "/path_dest/" PARALLELISM=5;

upvoted 1 times

musafir 2 years, 5 months ago

Selected Answer: ABC

ABC is correct, E is wrong since transporting DBs must have the same endian format.

upvoted 1 times

julmis 2 years, 5 months ago

Selected Answer: ABC

A, B , C

upvoted 1 times

Guhborges 2 years, 6 months ago

Selected Answer: BCE

B,C,E No doubt.

upvoted 1 times

kaka321 2 years, 7 months ago

I will go for ABC for me F is wrong as it is not automatically achived you need to run the covert command on the source db reference the target db. see below https://docs.oracle.com/en/database/oracle/oracle-database/18/bradv/rman-transporting-data-across-platforms.html#GUID-E681FF6F-0927-4ED5-BF64-73926C4121DC

upvoted 1 times

FelipeC 2 years, 7 months ago

Selected Answer: BC

E: incorrect If the destination database uses an endian format that is different from that of the source database, then the required endian format conversion is performed on the destination database (This is applicable for tablespace level transport) the question is about transporting databases. You can transport a entire database only if the source and destination use the same endian format. Doc ID 2013271.1

B, C correct for sure

upvoted 1 times

j_tw 2 years, 9 months ago

Selected Answer: ABC

E is incorrect. A tablespace can be transported between different endian format, but a Database should be the same endian format to be transported.

upvoted 1 times

ScottL 2 years, 9 months ago

Selected Answer: ABC

E incorrect. ABC correct

upvoted 1 times

ScottL 2 years, 9 months ago

Selected Answer: ABC

A,B,C

E is incorrect. From doc - "You can use RMAN to transport tablespaces across platforms with different endian formats. You can also use RMAN to transport an entire database to a different platform so long as the two platforms have the same endian format."

upvoted 1 times

gio 2 years, 10 months ago

Selected Answer: ABC

Not sure for A, I can't find anything in documentation about this. But E is surely incorrect: https://docs.oracle.com/en/database/oracle/oracle-database/19/bradv/rman-transporting-data-across-platforms.html#GUID-B0178538-6E41-455D-8166-FBB1EDC5D198

upvoted 1 times

gio 2 years, 6 months ago

why not D? I think D is better than A

upvoted 2 times

Question #18 *Topic 1*

Examine this command:
$ rhpctl move database `"sourcehome Oracle_home_path `"destinationhome Oracle_home_path
For which two purposes can you use this command? (Choose two.)

A. to switch an existing Oracle Database home to a newer release of Oracle software on the same server

B. to switch to a read-only Oracle home

C. to switch back to the previous Oracle home as part of a rollback operation

D. to switch the Oracle Database home when using a centralized Rapid Home Provisioning server

E. to switch to a patched Oracle Database home

Suggested Answer: *A, E*
Reference: https://docs.oracle.com/en/database/oracle/oracle-database/18/cwadd/rapid-home-provisioning.html#GUID-856C7378-2510-4B3E-8BE1- B4D79DF4FE8B

Jatindra Highly Voted 5 years, 8 months ago
Student Guide: You can use rhpctl move gihome command with the same syntax to switch from the current Oracle Grid Infrastructure home to a patched home. The rhpctl command enables you to switch from your current Oracle Grid Infrastructure or Oracle Database home to patched Oracle home so that you can provision the new Oracle home as gold image. You can also use the rhpctl command to switch back to the old Oracle home, if you want to roll back the operation. Correct Answer: CE option A is incorrect and confusing..
upvoted 17 times

mamadu Most Recent 1 year, 11 months ago
Answer = C,E True, A is incorrect. student guide: You can switch Oracle home and Oracle Grid infrastructure homes to later versions WITHOUT using a centralized Rapid Home Provisioning server.
upvoted 2 times

dancymonkey 2 years ago
E - Moves one or more databases from a source working copy or any Oracle Database home to a patched working copy.

https://docs.oracle.com/en/database/oracle/oracle-database/12.2/cwadd/rapid-home-provisioning-and-server-control-command-reference.html#GUID-C10036DC-B091-40E1-B46B-8ACF38B20066:~:text=Moves%20one%20or%20more%20databases%20from%20a%20source%20working%20copy%20or%20any%20Oracle%20Database%20home%20to%20a%20patched%20working%20copy.

C - think as backward of E

Answer is CE
upvoted 2 times

LEOC71 2 years, 2 months ago
Selected Answer: BE
A. False: if upgrading use "upgrade database" B. True: "move database" checks for all patches present in current home are also present in new home, so the new home can be identical to the old home at patches level but can be read only. C. False: to switch back use "-ignorewcpatches" D. False: "move database" doesn't requiere a centralized Rapid Home Provisioning server E. True: "move database" checks for all patches present in current home are also present in new home, but more patches can be present.
upvoted 2 times

greenday5494 2 years, 5 months ago
Selected Answer: DE
D and E seem correct. So does C..
upvoted 2 times

gio 2 years, 6 months ago
Selected Answer: CE
I think CE

upvoted 1 times

Alvinzzz 2 years, 8 months ago

B seems correct. If both source and destination Oracle Home are read-only, I can't see any additional actions required.

C seems incorrect. You can switch the home but should not consider as "rollback" operation. You have to specify "-rollback" in the command, else some previous configuration files and parameters may not restored as part of "rollback" operation.

upvoted 1 times

ScottL 2 years, 9 months ago

Selected Answer: CE

CE correct, see Jatindra

upvoted 1 times

vkra 3 years ago

Selected Answer: CE

see Jatindra

upvoted 1 times

hilaire 3 years, 2 months ago

C E correct

upvoted 1 times

arpz 4 years, 5 months ago

Answer should be C and E

upvoted 3 times

Neil107 4 years, 6 months ago

C and E are correct.

upvoted 2 times

Aldrid 4 years, 8 months ago

A. to switch an existing Oracle Database home to a newer release of Oracle software on the same server Incorrect: Must be same database release level.

B. to switch to a read-only Oracle home Incorrect

C. to switch back to the previous Oracle home as part of a rollback operation May be correct for rhpctl but not with "rhpctl move database" May be with —revert parametr

D. to switch the Oracle Database home when using a centralized Rapid Home Provisioning server Looks correct: This tool is Rapid Home Provisioning Control..so it's part of RHP Server..

E. to switch to a patched Oracle Database home Absolutely correct from command base description.

upvoted 4 times

Aldrid 4 years, 8 months ago

DE Use the rhpctl move database command to move one or more database homes to a working copy of the same database release level.

rhpctl is Rapid Home Provisioning Control.. So D is Correct

upvoted 3 times

Aldrid 4 years, 8 months ago

Moves one or more databases from a source working copy or any Oracle Database home to a patched working copy.

upvoted 1 times

ObserverPL 4 years, 9 months ago

Definitely C and E are ok

upvoted 1 times

Michael_14 5 years, 7 months ago

Why B is incorrect? C is not completely correct since rolling back patches requires the -ignorewcpatches parameter which is not mentioned in the question

upvoted 2 times

Jatindra 5 years, 8 months ago

Three options are correct.. ACE

upvoted 1 times

tamagogo 5 years, 8 months ago

Answer is CE

upvoted 1 times

Question #19 *Topic 1*

Which two are true about changing the LOCAL_UNDO_ENABLED property to false in a CDB? (Choose two.)

A. After the change, only a common user with the required privilege can create an undo tablespace in CDB&ROOT.

B. Any new PDB and existing PDBs are automatically configured to use the default undo tablespace in CDB$ROOT.

C. After the change, only one undo tablespace can exist in CDB$ROOT.

D. After the change, any user with the required privilege can create an undo tablespace in the PDBs.

E. Undo tablespaces existing in PDBs must be dropped before the change.

F. After the change, each existing PDB has to be reopened for the new undo mode to take effect.

Suggested Answer: *A, D*

Reference: https://docs.oracle.com/en/database/oracle/oracle-database/18/multi/creating-and-configuring-a-cdb.html#GUID-12ADA04D-F81D-4579-A68C- 0958CC7D6C2F

Community vote distribution

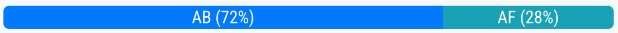

ama Highly Voted 5 years, 7 months ago

A is correct

B is wrong, because PDBs must be reopened to recognize the central UNDO

C is wrong because as the database is RAC and enabled with LOCAL UNDO, you need to assign specific Undo tablespace to each PDB on node-1 & node-2 by executing below commands. On node-1: alter system set undo_tablespace=UNDOTBS1 container=current sid='DBCDB1' scope=spfile; >>>> sid=First Instance On Node-2 alter system set undo_tablespace=UNDOTBS2 container=current sid='DBCDB2' scope=spfile; >>>> sid=Second Instance

D is wrong becuase as long as LOCAL UNDO used you cant create UNDO in PDB

E is wrong

F is correct

upvoted 23 times

dasigrist 5 years, 5 months ago

C is also wrong for a single instance too because after you change from true to false, the existing PDB's undo tablespace will still be there. They are not deleted automatically. They are not going to be used, but they are still there.

upvoted 4 times

Franky_T Highly Voted 4 years, 3 months ago

Selected Answer: AB

A is correct. In shared UNDO mode, UNDO can exist in the CDB only. B is correct. The change is automatic and comes into effect as soon as you restart the CDB instance (required). C is wrong. More than one UNDO tablespace can exist in the CDB, but only one can be the actice UNDO tablespace. D is wrong. Tested this. Switched to Shared Undo, connected to PDB as SYS, created an undo tablespace (create statement succeeded). Query dba_tablespaces and dba_data_files from CDB and the tablespace nor the datafile exists. Conslusion - create undo tablespace command ignored. E is wrong. PDB UNDO tablespaces should be dropped after the change. F is wrong. The CDB has to be restarted after the change so no need to reopen each PDB.

upvoted 10 times

zenai Most Recent 1 year, 9 months ago

Selected Answer: AF

A, F https://docs.oracle.com/en/database/oracle/oracle-database/19/multi/administering-a-cdb-with-sql-plus.html#GUID-24EA5811-94F0-4EEC-864F-23AEF48F2D51

upvoted 1 times

auwia 2 years, 1 month ago

Selected Answer: AF

See ama's comment

upvoted 1 times

Kuraudio 2 years, 6 months ago

Selected Answer: AB

According to this: https://docs.oracle.com/en/database/oracle/oracle-database/19/multi/administering-a-cdb-with-sql-plus.html#GUID-D54AAD88-03DA-4E22-991F-8C15306823A5

upvoted 1 times

dancymonkey 2 years, 6 months ago

AF

F - PDB has to be reopened as part of startup upgrade in CDB

ALTER DATABASE LOCAL UNDO OFF - ORA-65192: database must be in UPGRADE mode for this operation

SQL> shutdown immediate; ..... SQL> startup upgrade; ORACLE instance started.

Total System Global Area 1124073328 bytes Fixed Size 9133936 bytes Variable Size 738197504 bytes Database Buffers 369098752 bytes Redo Buffers 7643136 bytes Database mounted. Database opened. SQL> ALTER DATABASE LOCAL UNDO OFF;

Database altered.

SQL> show pdbs;

CON_ID CON_NAME OPEN MODE RESTRICTED ---------- ------------------------------ ---------- ---------- 2 PDB$SEED MIGRATE YES 3 ORCLPDB MOUNTED SQL>

upvoted 1 times

LEOC71 2 years, 8 months ago

Selected Answer: AB

C. False. Only one active, but many created. D. False. Command to create undo tablespaces are ignored with no error. E. False. Drop is recommended after. F. False. The CDB is restarted as part of the procedure, not individual PDBs.

upvoted 1 times

musafir 2 years, 11 months ago

Selected Answer: AF

A and F appear to be correct https://docs.oracle.com/en/database/oracle/oracle-database/19/multi/administering-a-cdb-with-sql-plus.html#GUID-D54AAD88-03DA-4E22-991F-8C15306823A5

upvoted 1 times

lchdb 3 years ago

i want AB C : false - When in shared undo mode, the CDB ignores any local undo tablespaces that were created when it was in local undo mode. Oracle recommends that you delete the unused local undo tablespaces. (only=x / can exist more) D : false - not (in the PDBs) / -> in CDB$ROOT E : false - reference C / not must F : false - The CDB should be opened, not the PDB. But from PDB's point of view, F might be right. In terms of (in a CDB), I think A and B are better

upvoted 1 times

ScottL 3 years, 3 months ago

Selected Answer: AB

AB I think

upvoted 1 times

gio 3 years, 4 months ago

Selected Answer: AF

B is wrong because you need to restart CDB: https://docs.oracle.com/en/database/oracle/oracle-database/19/multi/administering-a-cdb-with-sql-plus.html#GUID-D54AAD88-03DA-4E22-991F-8C15306823A5

upvoted 2 times

smartvan 4 years, 9 months ago

I think is AF

A - Only a common user who has the appropriate privileges and whose current container is the CDB root can create an undo tablespace.

F - When you change the undo mode of a CDB, the new undo mode applies to an individual container the first time the container is opened after the change.

C is incorrect - When it is in shared undo mode, the CDB ignores any local undo tablespaces that were created when it was in local undo mode. Oracle recommends that you delete these local undo tablespaces.

So Oracle RECOMMENDS, is not mandatory.

https://docs.oracle.com/en/database/oracle/oracle-database/12.2/admin/administering-a-cdb-with-sql-plus.html#GUID-1AAB032F-D13A-4E06-A062-59E0487C2FC5

upvoted 1 times

chefdba 4 years, 9 months ago

AF = correct

upvoted 1 times

senator 4 years, 11 months ago

A & F are the correct answers

upvoted 1 times

Neil107 5 years ago

I think it's BC. You can set a CDB in local UNDO mode either at CDB creation or by altering the CDB property. When the database property LOCAL_UNDO_ENABLE is FALSE, which is the default, there is only one UNDO tablespace that is created in the CDB root, and that is shared by all containers. When LOCAL_UNDO_ENABLE is TRUE, every container in the CDB uses local undo and each PDB must have its own local UNDO tablespace. To maintain ease of management and provisioning, UNDO tablespace creation happens automatically and does not require any action from the user. When a PDB is opened and an UNDO tablespace is not available, its automatically created.

upvoted 1 times

wiprooracle 5 years, 2 months ago

A, C looks correct as per

When a CDB is in shared undo mode, the following applies:

Only a common user who has the appropriate privileges and whose current container is the CDB root can create an undo tablespace.

When the current container is not the CDB root, an attempt to create an undo tablespace fails and returns an error.

From https://docs.oracle.com/en/database/oracle/oracle-database/12.2/admin/administering-a-cdb-with-sql-plus.html#GUID-8F8B2FF8-7FA7-40CD-8AA5-ACABCD3964D8

upvoted 1 times

ABAJ 5 years, 2 months ago

BF are correct

Shared Undo Mode

Shared undo mode means that there is one active undo tablespace for a single-instance CDB. For an Oracle RAC CDB, there is one active undo tablespace for each instance.

When a CDB is in shared undo mode, the following applies:

Only a common user who has the appropriate privileges and whose current container is the CDB root can create an undo tablespace.

When the current container is not the CDB root, an attempt to create an undo tablespace fails and returns an error.

Undo tablespaces are visible in static data dictionary views and dynamic performance (V$) views when the current container is the CDB root. Undo tablespaces are visible only in dynamic performance views when the current container is a PDB, an application root, or an application PDB.

upvoted 1 times

ABAJ 5 years, 1 month ago

AF are correct.

upvoted 2 times

Question #20 *Topic 1*

Which two are true about SQL Performance Analyzer (SPA)? (Choose two.)

A. It is integrated with the SQL Access Advisor.

B. It predicts the impact of system changes on SQL workload response time.

C. It provides before and after execution statistics for each SQL statement in the analysis task

D. It offers fine-grained analysis of all the SQL statements in the analysis task as a group.

E. SQL statements that were originally run concurrently are run concurrently by SPA.

Suggested Answer: *B, D*

Reference: https://www.oracle.com/technetwork/database/manageability/spa-ow09-131455.pdf (4)

Community vote distribution



dacoben415lywenw Highly Voted 5 years, 1 month ago

I think BC

B: Oracle Real Application Testing ... enables businesses to fully assess the outcome of a system changes in test or production. Oracle Real Application Testing is comprised of the following features: Performance Analyzer (SPA), SPA ... C: SQL Performance Analyzer executes the SQL statements captured in the SQL Tuning Set and generates execution plans and execution statistics for each statement

For D , though group (summary) view is shown during analysis, i dont think that what they mean by fine grained

upvoted 20 times

kaka321 Most Recent 2 years, 1 month ago

D is wrong because it provides details per query and summary as a group.

upvoted 1 times

ScottL 2 years, 3 months ago

Selected Answer: BC

B C correct

upvoted 1 times

gio 2 years, 4 months ago

Selected Answer: BC

I think BC is correct

upvoted 1 times

vkra 2 years, 6 months ago

Selected Answer: BC

see dacoben415lywenw

upvoted 1 times

DBA4All 3 years, 2 months ago

B & C: https://docs.oracle.com/en/database/oracle/oracle-database/19/ratug/introduction-to-sql-performance-analyzer.html#GUID-F07105CA-0DFA-4102-A8FE-1C11AD201528

upvoted 4 times

kameni 3 years, 12 months ago

Answer is B and C B --> Helps users predict the impact of system changes on SQL workload response time C--> SQL Performance Analyzer compares the performance of SQL statements before and after the change and produces a report identifying any changes in execution plans or performance of the SQL statements.

upvoted 3 times

avanand 4 years, 11 months ago

I will go with B D.

SPA offers fine grained analysis for each SQL. See the doc below https://www.oracle.com/technetwork/database/manageability/spa-ow08-131564.pdf

upvoted 2 times

mihanik 4 years, 10 months ago

*But not as a group B: SQL Performance Analyzer measures the impact of system changes both on the overall execution time of the SQL workload and on the response time of every individual SQL statement in the workload. By default, SQL Performance Analyzer uses elapsed time as a metric for comparison.

I vote for BC

upvoted 4 times

cerebro2000x 5 years ago

BC based on my experience

upvoted 4 times

veginha 5 years, 2 months ago

B is correct. i'm not pretty sure if D is correct, as i understand that fine-grained analsys is offered by individual SQL not as group. Anyone? https://www.oracle.com/a/otn/docs/enterprise-manager/ds-19c-oracle-real-application-testing.pdf

upvoted 1 times

Question #21 *Topic 1*

A user complains about poor database performance.
You want to verify if the user's session has waited for certain types of I/O activity.
Which view displays all waits waited on by a session at least once?

A. V$SESSION_EVENT

B. V$SESSTAT

C. V$SESSION_WAIT

D. V$SESSION_WAIT_CLASS

E. V$SESSION

Suggested Answer: *D*
Reference: https://docs.oracle.com/cd/E29597_01/server.1111/e16638/instance_tune.htm

Community vote distribution
A (100%)

janw Highly Voted 4 years, 8 months ago
V$SESSION_EVENT This view lists information on waits for an event by a session.
upvoted 13 times

janw Highly Voted 4 years, 8 months ago
correct: A
upvoted 7 times

gio Most Recent 1 year, 10 months ago
Selected Answer: A
I think A
upvoted 1 times

ErikJanssen 2 years, 3 months ago
A and D I think, are correct
upvoted 1 times

Patrick9230 2 years, 7 months ago
Selected Answer: A
must be A
upvoted 3 times

Dhanushbh 3 years, 7 months ago
V$SESSION_EVENT The V$SESSION_EVENT view provides summary of all the events the session has waited for since it started.
upvoted 1 times

navingupta52 3 years, 8 months ago
Option A is correct.

https://docs.oracle.com/cd/E29597_01/server.1111/e16638/instance_tune.htm
upvoted 3 times

flaviogcmelo 3 years, 10 months ago
The most correct answer is A. V$SESSION_EVENT, but C. V$SESSION_WAIT can displays the current or last wait for each session.
https://docs.oracle.com/en/database/oracle/oracle-database/19/refrn/V-SESSION_EVENT.html#GUID-5723D404-83F7-4C8D-8BF2-031B3FFEAAA8
https://docs.oracle.com/en/database/oracle/oracle-database/19/refrn/V-SESSION_WAIT.html#GUID-4EDAB293-F3FC-40FE-BC75-4FEE6A4D7705
upvoted 2 times

cerebro2000x 4 years, 6 months ago
V$SESSION_EVENT This view lists information on waits for an event by a session
upvoted 4 times

Question #22 *Topic 1*

Which two are true about gathering optimizer statistics? (Choose two.)

A. Executing DBMS_STATS.GATHER_DATABASE_STATS while connected to CDB$ROOT gathers object statistics in all open PDBs except PDB$SEED.

B. Executing DBMS_STATS.GATHER_DATABASE_STATS while connected to a PDB opened in read/write mode gathers object statistics for that PDB.

C. Executing DBMS_STATS.GATHER_DATABASE_STATS while connected to CDB$ROOT gathers object statistics only in CDB$ROOT.

D. System statistics can be gathered only while connected to CDB$ROOT.

E. Executing DBMS_STATS.GATHER_DATABASE_STATS while connected to CDB$ROOT gathers object statistics in all open pluggable databases (PDBs)

Suggested Answer: *C, E*

Reference: https://docs.oracle.com/en/database/oracle/oracle-database/20/arpls/DBMS_STATS.html

Community vote distribution

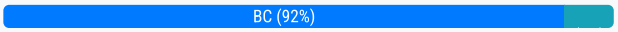

antoniomariano1983 Highly Voted 4 years, 2 months ago

Selected Answer: BC

BC test: SQL> select count(*) from sys.dba_tab_statistics where last_analyzed is not null; COUNT(*) --------- 3610 SQL> EXEC DBMS_STATS.GATHER_DATABASE_STATS; PL/SQL procedure successfully completed. SQL> select count(*) from sys.dba_tab_statistics where last_analyzed is not null; COUNT(*) --------- 3649 SQL> alter session set container=ORCLPDB1; Session altered. SQL> select count(*) from sys.dba_tab_statistics where last_analyzed is not null; COUNT(*) --------- 3507 SQL> EXEC DBMS_STATS.GATHER_DATABASE_STATS; PL/SQL procedure successfully completed. SQL> select count(*) from sys.dba_tab_statistics where last_analyzed is not null; COUNT(*) --------- 3598

upvoted 7 times

ABAJ Highly Voted 5 years, 3 months ago

BC are correct. Tested in LAB.

upvoted 6 times

fiko666 4 years, 7 months ago

How can you test that?

upvoted 2 times

cujar2003 Most Recent 1 year, 9 months ago

Selected Answer: BE

B y E Test in my LAB 19C SQL> select max(last_analyzed) from sys.dba_tab_statistics;

MAX(LAST_ANALYZED) ---------- 14-OCT-24

SQL> show pdbs

CON_ID CON_NAME OPEN MODE RESTRICTED --------- ---------- ---------- --------- 2 PDB$SEED READ ONLY NO 3 PDBDEV READ WRITE NO SQL> show con_name

CON_NAME ---------- CDB$ROOT SQL> EXEC DBMS_STATS.GATHER_DATABASE_STATS;

PL/SQL procedure successfully completed.

SQL> select max(last_analyzed) from sys.dba_tab_statistics;

MAX(LAST_ANALYZED) ---------- 04-NOV-24

SQL> alter session set container=PDBDEV;

Session altered.

SQL> select max(last_analyzed) from sys.dba_tab_statistics;

MAX(LAST_ANALYZED) ---------- 04-NOV-24

SQL>

upvoted 1 times

Guhborges 3 years, 1 month ago

Selected Answer: BC

B and C are correct.

upvoted 1 times

Alvinzzz 3 years, 2 months ago

For option C, the tricky point is the word "only. When connected to CDB$ROOT, you can also collect object statistics at PDB level if specified.

For option D, all containers under the CDB share the same hardware. I can't see why it was incorrect as DBMS_STATS.GATHER_SYSTEM_STATS should be executed at CDB level and applies to the PDBs under it.

upvoted 1 times

ScottL 3 years, 3 months ago

Selected Answer: BC

BC correct

upvoted 1 times

gio 3 years, 4 months ago

Selected Answer: BC

I think BC

upvoted 1 times

jareach 3 years, 6 months ago

Selected Answer: BC

Suggested CE are mutuallz exclusive.

upvoted 2 times

Kushal22 3 years, 7 months ago

BC is correct .https://docs.oracle.com/en/database/oracle/oracle-database/21/arpls/DBMS_STATS.html#GUID-9761DFB9-8710-45D9-9ACE-A4FD37FD69F5

The application PDB user, uses this preference to determine whether to allow the application root to interact with the statics gathering in PDB.

During the statistics gathering of a metadata linked table in the application root, if the statistics in a PDB are in stale state, the application root triggers the statistics gathering for the particular PDB. Using this preference, the user can either execute or ignore the command from the application root. You can set the following values:

TRUE—Allows the application root trigger the statistics gathering on metadata linked table in application PDB if the statistics on PDB are stale.

FALSE—Ignores the statistics gathering command triggered from application root.

The default value is FALSE.

Note:

CDB root, different from application root, never triggers statistics gathering on the PDBs and it is not controlled by this preference.

upvoted 1 times

chefdba 4 years, 9 months ago

BC = 100% correct

upvoted 2 times

rashedptdba 4 years, 10 months ago

B,C are correct answer. Tested on my lab

upvoted 2 times

Neil107 5 years ago

B and E

upvoted 2 times

navingupta52 5 years, 1 month ago

B & C are correct. Tested on my system.

upvoted 3 times

ObserverPL 5 years, 3 months ago

B is obvious, also C is true as I've just tested in practice on 19c - EXEC DBMS_STATS.GATHER_DATABASE_STATS launched on CDB$ROOT did NOT calculated stats on PDB1:EMP table I created.

upvoted 3 times

cerebro2000x 6 years ago

SQL> show pdbs

CON_ID CON_NAME OPEN MODE RESTRICTED --------- ---------- ---------- 2 PDB$SEED READ ONLY NO 3 HR READ WRITE NO SQL> alter session set container = PDB$SEED;

Session altered.

SQL> show con_name

CON_NAME ---------- PDB$SEED SQL> select count(*) from sys.dba_tab_statistics where last_analyzed is not null;

COUNT(*) ---------- 2290

SQL> EXEC DBMS_STATS.GATHER_DATABASE_STATS; BEGIN DBMS_STATS.GATHER_DATABASE_STATS; END;

* ERROR at line 1: ORA-20000: database or pluggable database open for read-only access ORA-06512: at "SYS.DBMS_STATS", line 42559 ORA-06512: at "SYS.DBMS_STATS", line 42521 ORA-06512: at "SYS.DBMS_STATS_INTERNAL", line 17263 ORA-06512: at "SYS.DBMS_STATS", line 42414 ORA-06512: at "SYS.DBMS_STATS", line 42549 ORA-06512: at line 1

upvoted 2 times

janw 6 years ago

BC:

https://mikedietrichde.com/2016/10/21/gather-fixed-objects-stats-in-pdbs-as-well/

upvoted 2 times

dacoben415lywenw 6 years, 1 month ago

corect BC tested on 19c and 12.2 test provided by ald85 is not accurate, E is not corect try creating simple tales in ROOT and PDB than gather db stats from ROOT, and you will see that PDB tables stats will not be gathered.

upvoted 3 times

Question #23 *Topic 1*

Examine this output:

Which two are true? (Choose two.)

```
SQL> select pluggable_database, shares, parallel_server_limit
  2 from dba_cdb_rsrc_plan_directives where plan = 'MY_PLAN'
  3 order by plugaable_database;

PLUGGABLE_DATABASE            SHARES    PARALLEL_SERVER_LIMIT
----------------------------- --------- ----------------------------
ORA$AUTOTASK                                                    100
ORA$DEFAULT_PDB_DIRECTIVE             1                           0
PDB1                                  2                         100
PDB2                                  2                          25
PDB3                                  1
SQL> select name, value from v$parameter
  2 where name = 'resource_manager_plan';

NAME                      VALUE
------------------------- ------------
resource_manager_plan     MY_PLAN
```

A. Any PDB not specified in the plan will be unable to execute statements in parallel.

B. PDB3 can use all available parallel execution processes at times.

C. PDB1 is always limited to 40% of the available system resources regardless of demand.

D. Any PDB not specified in the plan will be able to use a maximum of 16.5% of the available system resources.

E. PDB3 is guaranteed to receive at least 20% of the available system resources if there is enough demand.

F. PDB2 is guaranteed at least 25% of the available parallel execution processes if there is enough demand.

Suggested Answer: *A, E*

Official Answer: AE

Community vote distribution

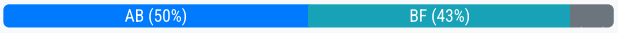

SMagny Highly Voted 5 years, 8 months ago

A and B for me. A) PDB not specified in the Plan will be assigne to the DEFAULT_PDB_DIRECTIVE. This directive set PARALLEL_SERVER_LIMIT to zero. (No parallel server). B) PARALLEL_SERVER_LIMIT is not set for PDB3 Resource plan directive. The default value of 100% is used in this case.

upvoted 24 times

Kuraudio 2 years, 7 months ago

The default value is 100 as described here: https://docs.oracle.com/en/database/oracle/oracle-database/18/multi/using-oracle-resource-manager-for-pdbs-with-sql-plus.html#GUID-3E84009E-5CD6-46B8-84D3-BE775D58D9AA

upvoted 1 times

Alejandrrro Highly Voted 5 years, 6 months ago

Agree with you

A. Correct as then a PDB will get assigned the default directive with parallel_server_limit set to 0 B. Correct as parallel_server_limit is not specified for PDB3 and it defaults to 100

Quote:

https://docs.oracle.com/database/121/ADMIN/dbrm.htm#ADMIN14008

"If PARALLEL_SERVER_LIMIT is not specified, then it defaults to 100%."

C. Incorrect. A PDB can use more resources:

Quote from

https://docs.oracle.com/database/121/ADMIN/cdb_dbrm.htm#ADMIN13777

"However, any PDB can use more than the guaranteed amount of a resource if there is no resource contention."

D. Incorrect. Shares is a minimum guaranteed value. But a PDB can use more as explained above in C. E. Incorrect. Nothing confirms that there are only 3 PDBs attached. There may be others using the default pdb directive. Then the share for PDB2 is not 20% F. Incorrect. parallel_server_limit is an upper limit.

upvoted 5 times

asduta 3 years, 7 months ago

I was wondering that since the default directive with parallel_server_limit set to 0, how is it that other PDB will get share (for your explaination of F) since they wont be allowed to run parallel?

upvoted 1 times

cujar2003 Most Recent 1 year, 9 months ago

Selected Answer: BF

B y F i agree SQL> COLUMN plan FORMAT A30 COLUMN pluggable_database FORMAT A25 SET LINESIZE 100

SQL> SQL> SQL> SQL> SELECT plan, pluggable_database, shares, utilization_limit AS util, parallel_server_limit AS parallel FROM dba_cdb_rsrc_plan_directives WHERE plan = 'TEST_CDB_PLAN' ORDER BY pluggable_database; 2 3 4 5 6 7 8

PLAN PLUGGABLE_DATABASE SHARES UTIL PARALLEL ---------- ---------- ---------- ---------- ---------- TEST_CDB_PLAN ORA$AUTOTASK 90 TEST_CDB_PLAN ORA$DEFAULT_PDB_DIRECTIVE 1 TEST_CDB_PLAN PDBDEV 2 TEST_CDB_PLAN PDBDEV2 2 25 TEST_CDB_PLAN PDBDEV3 1

upvoted 1 times

cujar2003 1 year, 9 months ago

Selected Answer: AB

SQL> SQL> SQL> SQL> SELECT plan, pluggable_database, shares, utilization_limit AS util, parallel_server_limit AS parallel FROM dba_cdb_rsrc_plan_directives WHERE plan = 'TEST_CDB_PLAN' ORDER BY pluggable_database; 2 3 4 5 6 7 8

PLAN PLUGGABLE_DATABASE SHARES UTIL PARALLEL ---------- ---------- ---------- ---------- ---------- TEST_CDB_PLAN ORA$AUTOTASK 90 TEST_CDB_PLAN ORA$DEFAULT_PDB_DIRECTIVE 1 TEST_CDB_PLAN PDBDEV 2 TEST_CDB_PLAN PDBDEV2 2 25 TEST_CDB_PLAN PDBDEV3 1

SQL>

upvoted 1 times

Soiram 1 year, 10 months ago

B is correct no doubt. For me the second correct is E. One share of five is 20%. Also we can conclude that these are all the PDBs, since there is not another reciurce_manager_plan.

upvoted 1 times

Soiram 1 year, 10 months ago

Sorry, E is ont correct as it is NOT guaranteed. A is correct because default plan is 0

upvoted 1 times

auwia 2 years, 1 month ago

Selected Answer: AB

A and B.

upvoted 1 times

Guhborges 3 years, 1 month ago

Selected Answer: AB

A,B for sure

upvoted 1 times

gio 3 years, 4 months ago

Selected Answer: AB

I think AB

upvoted 1 times

Franta1 3 years, 10 months ago

Selected Answer: AB

Agree with SMagny and Alejandrrro. Options A,B

upvoted 3 times

flaviogcmelo 4 years, 3 months ago

Selected Answer: AE

When you do not explicitly define directives for a PDB, the PDB uses the default directive for PDBs. In this case the default directive sets the value 0 for parallel_server_limit. In this case I understand that PDB3 cannot use resources in parallel. So option A seems correct to me.

upvoted 1 times

flaviogcmelo 4 years, 3 months ago

The correct options ar A and F in my opinion.

upvoted 1 times

mtnetmaker 4 years, 3 months ago

A- Incorrect : If PARALLEL_SERVER_LIMIT is not specified, then it defaults to 100% REF: https://docs.oracle.com/database/121/ADMIN/dbrm.htm#ADMIN14007 https://docs.oracle.com/en/database/oracle/oracle-database/19/multi/using-oracle-resource-manager-for-pdbs-with-sql-plus.html#GUID-3E84009E-5CD6-46B8-84D3-BE775D58D9AA When you do not explicitly define directives for a PDB, the PDB uses the default directive for PDBs. B - Correct : If PARALLEL_SERVER_LIMIT is not specified, then it defaults to 100% C - Incorrect: PDB1 is NOT always limited to 40% of the available system resources regardless of demand. D - Incorrect : When you do not explicitly define directives for a PDB, the PDB uses the default directive for PDBs., default is 100% E: Incorrect : available system resources = utilization_limit and parallel_server_limit F: Correct : parallel_server_limit =25

upvoted 2 times

asduta 4 years, 2 months ago

A is correct since itll default to DEFAULT_PDB_DIRECTIVE that was set to 0.

upvoted 1 times

vkra 3 years, 6 months ago

F: is incorrect, "parallel_server_limit" does not guarantee anything, it limits the resource to 25% !
https://docs.oracle.com/database/121/ADMIN/cdb_dbrm.htm#ADMIN13786

upvoted 1 times

casfdsaf 4 years, 4 months ago

Selected Answer: BF

ithink bf

upvoted 2 times

yarsalan 4 years, 6 months ago

Selected Answer: BF

In my opinion, B,F are correct.

upvoted 3 times

flaviogcmelo 4 years, 9 months ago

Question:

Which three are true about performing an Oracle Database install on Linux? (Choose three.)

Response:

A. The Oracle database administrator must be granted access to the root operating system account to tun root privileged scripts. B. It allows you to select the languages supported by the Oracle database server. C. The runfixup.sh script can install missing RPMs. D. The Oracle Preinstallation RPM can be used to configure the Oracle database installation owner, the Oracle Inventory group, and an Oracle administrative privileges group. E. It can be done after installing Grid Infrastructure for a Standalone Server. F. It can be done before installing Grid Infrastructure for a Standalone Server. G. The Oracle Preinstallation RPM must be used to configure the Oracle database installation owner, the Oracle Inventory group, and an Oracle administrative privileges group.

D, E and F are correct. Any comments???

upvoted 1 times

chefdba 4 years, 9 months ago

BF = correct answers

upvoted 1 times

Neil107 5 years ago

I believe it's AE.

upvoted 3 times

xsxxxxx 5 years, 2 months ago

not A, "Any PDB that is part of the CDB and does not have directives defined for it uses the default PDB directive. "

upvoted 2 times

asduta 3 years, 7 months ago

and DEFAULT_PDB_DIRECTIVE parallel is 0. so, A is correct.

upvoted 1 times

Glykon 5 years, 7 months ago

A and E E because of: https://blog.dbi-services.com/cdb-resource-plan-shares-and-utilization_limit/

upvoted 2 times

Question #24 *Topic 1*

Which two are true about the execution of operating system scripts starting from Oracle Database 19c? (Choose two.)

A. orainstRoot.sh can be executed automatically by the Database installer by using sudo or root credentials.

B. root.sh can be executed automatically by the Database Installer only if it is provided with root credentials.

C. The sudo password can be specified in a response file.

D. root.sh can be executed automatically by the Database installer only by using sudo credentials.

E. The sudo password must be specified in a response file.

F. The root password cannot be specified in a response file.

Suggested Answer: *A, F*
Official Answer: AF

Community vote distribution
AF (100%)

Jatindra Highly Voted 5 years, 8 months ago
Correct Answer :A,F
upvoted 13 times

Amaora 4 years, 8 months ago
F ???? I think you are wrong, look at this parameter ROOT_ACCESS_PASSWORD. This parameter enables you to input root's password in response file
upvoted 2 times

Kishor123 4 years, 7 months ago
I think you are referring to this document , this is not related to oracle s/w installation
https://docs.oracle.com/cd/E89154_01/STAIS/silentmode002.htm#STAIS1938

Correct Answer : AF

You can not mention root password in response file

** With or without response file(put all arguments in the command line $ORACLE_HOME/runInstaller -silent ....) you have to input root password when prompted ** using OUI - You can mention either enter root password or sudo file in installation page You continue to have the option to run the root configuration scripts manually.
upvoted 2 times

Ligarius Highly Voted 4 years, 8 months ago
Sorry guys... Correct : A-C
upvoted 8 times

mamadu Most Recent 1 year, 11 months ago
No password can be specified in the response file rather the installer prompts for the password on the console. So C is wrong.
upvoted 1 times

niabs 4 years, 7 months ago
I also think A and C
upvoted 2 times

Yennism 2 years, 2 months ago
AE check this link: We can write root password (encryted) in response file
https://docs.oracle.com/cd/E89154_01/STAIS/silentmode002.htm#STAIS1636
upvoted 2 times

ScottL 2 years, 9 months ago
Selected Answer: AF
A,F correct see Franky_T
upvoted 1 times

gio 2 years, 10 months ago

Selected Answer: AF

https://docs.oracle.com/en/database/oracle/oracle-database/19/tutorial-install-oracle-database-with-automatic-root-scripts-execution/files/db.rsp

upvoted 1 times

Patrick9230 3 years, 7 months ago

Selected Answer: AF

A F is correct

upvoted 2 times

Franky_T 3 years, 8 months ago

A is correct. OUI step 6 of 10 provides you with this choice. Can be either. B is wrong. Can be either using the supplied root password or the sudo credentials. C is wrong. See point E. D is wrong. Can be either using the supplied root password or the sudo credentials. E is wrong. The root or sudo password cannot be specified in the response file. F is correct. See point E.

upvoted 4 times

smartvan 4 years, 3 months ago

A correct - If you have performed a Management Agent installation using the agentDownload script, then run the root.sh script. Also, run the oraInstroot.sh script if this is the first Oracle product on the host. To do so: Log in as a root user on the host. Alternatively, use the sudo command to change to a root user.

I believe F is correct - there is nowhere mentioned that you need to add root password in response file

upvoted 2 times

nobody347 4 years, 4 months ago

C is right When you modify a response file template and save a file for use, the response file may contain plain text passwords. Ownership of the response file should be given to the Oracle software installation owner only, and permissions on the response file should be changed to 600. Oracle strongly recommends that database administrators or other administrators delete or secure response files when they are not in use.

upvoted 1 times

vska 3 years, 10 months ago

there are plaintext password for SYS, SYSTEM but not root or sudo passwords. At least in v.19, just try by yourself runInstaller. It shows all possible variables. So AF

upvoted 1 times

Question #25 *Topic 1*

Automatic Shared Memory Management is disabled for one of your database instances.
Some SQL statements perform poorly due to excessive hard parse activity, thereby degrading performance.
What would be your next step?

A. Run the SQL Access Advisor.

B. Run the Memory Advisor for the shared pool.

C. Run the SQL Tuning Advisor.

D. Run the Memory Advisor for the Program Global Area.

E. Run the Memory Advisor for the System Global Area.

Suggested Answer: *B*
Official Answer: B

Community vote distribution
B (75%) C (25%)

Jatindra Highly Voted 5 years, 8 months ago
B is correct
upvoted 19 times

Alejandrrro 5 years ago
Agree with you:

https://blogs.oracle.com/oraclemagazine/using-database-advisors

The shared pool advisor provides information about the estimated parse time in the shared pool for different pool sizes, the buffer cache advisor provides information about physical reads and time for the cache size, and the PGA advisor provides information about cache hit percentage against PGA target memory size.
upvoted 3 times

Franky_T Highly Voted 3 years, 8 months ago
Memory advisors are available under the following conditions: AMM - Memory Advisor, ASMM - SGA Advisor & PGA Advisor, If MSMM (manual shared memory management) & Automatic PGA Memory Management - Buffer Cache Advisor & PGA Advisor. And, APMM is the default for Oracle databases.

A is wrong. SQL Access Advisor is normally run after the SQL Tuning Advisor, if the latter does not give the required results. B is wrong. You cannot run the memory advisor for the shared pool only. C is correct. This is the logical first step in tuning SQL statements when no memory advisors are available to use. D is wrong. Hard parsing indicates that the statement in not cached in the shared pool, part of the SGA, and has nothing to do with the PGA. E is wrong. The SGA advisor is only available when running in ASMM mode.
upvoted 6 times

dancymonkey Most Recent 2 years ago
C is correct.

https://docs.oracle.com/en/database/oracle/oracle-database/21/tgsql/sql-tuning-guide.pdf

1.5.1 SQL Tuning Tasks 4. Defining the scope of the problem The scope of the solution must match the scope of the problem. Consider a problem at the database level and a problem at the statement level. For example, the shared pool is too small, which causes cursors to age out quickly, which in turn causes many hard parses. Using an initialization parameter to increase the shared pool size fixes the problem at the database level and improves performance for all sessions. However, if a single SQL statement is not using a helpful index, then changing the optimizer initialization parameters for the entire database could harm overall performance. If a single SQL statement has a problem, then an appropriately scoped solution addresses just this problem with this statement.
upvoted 2 times

guretto 2 years, 2 months ago

c "excessive hard parse" means that there is bad SQL (i.e, use of literals and no bind variables when the statement structure is the same)

upvoted 1 times

LEOC71 2 years, 2 months ago

Selected Answer: B

v$shared_pool_advice

upvoted 1 times

musafir 2 years, 5 months ago

Selected Answer: B

B is correct

upvoted 1 times

Guhborges 2 years, 6 months ago

Selected Answer: B

B is correct, it advises on memory size.

upvoted 1 times

kaka321 2 years, 7 months ago

B is correct it advises on proper memory size for manual memory configuration when AMM is not in place. see memory architecture for oracle https://docs.oracle.com/en/database/oracle/oracle-database/19/admin/managing-memory.html#GUID-24312D78-522D-4C08-8934-820BE18D34C7

upvoted 1 times

gio 2 years, 10 months ago

Selected Answer: C

I think C because Memory advisor doesn't exists for shared pool and SQL Tuning Advisor is the only option when Manual Shared Memory management is enabled

upvoted 1 times

hilaire 3 years, 5 months ago

B is correct

upvoted 1 times

erial 3 years, 10 months ago

B,D and E are false since ASMM is turned off and we cannot run memory advisors, the correct answer is C sql tuning advisor because hard parse is due either shared pool is too small or sql statements are not using bind variables.

upvoted 1 times

nobody347 4 years, 4 months ago

B. A hard parse rate of more than 100 a second indicates that there is a very high amount of hard parsing on the system. High hard parse rates cause serious performance issues and must be investigated. Usually, a high hard parse rate is accompanied by latch contention on the shared pool and library cache latches

upvoted 2 times

Dhanushbh 4 years, 7 months ago

B is correct , no doubt about it

upvoted 3 times

asefa 5 years ago

i will go with B.

upvoted 4 times

klever35 5 years, 8 months ago

I still have the doubt. Run the Memory Advisor for the Program Global Area

upvoted 2 times

Question #26 *Topic 1*

Which two are true about flashback features in Oracle Database 19c and later releases? (Choose two.)

A. Flashback logs are automatically purged when DB_FLASHBACK_RETENTION_TARGET is set lower than the time they have already been retained.

B. Flashback logs are monitored and proactively deleted when beyond the retention period defined in DB_FLASHBACK_RETENTION_TARGET only after there is space pressure.

C. Flashback logs are monitored and proactively deleted when beyond the retention period defined in DB_FLASHBACK_RETENTION_TARGET before there is space pressure.

D. Flashback logs are monitored for being older than the retention period defined in DB_FLASHBACK_RETENTION_TARGET and can be deleted by an administrator written event trigger.

E. Flashback logs are automatically purged whenever the value of DB_FLASHBACK_RETENTION_TARGET is changed.

Suggested Answer: *A, C*
Official Answer: AC

Community vote distribution



Adel_007 Highly Voted 5 years, 11 months ago
Starting with Oracle Database Release 19c, the management of space in the fast recovery area is simplified. Oracle Database monitors flashback logs in the fast recovery area and automatically deletes flashback logs that are beyond the retention period. When the retention target is reduced, flashback logs that are beyond the retention period are deleted immediately.

In scenarios where a sudden workload spike causes a large number of flashback logs to be created, the workload is monitored for a few days before deleting flashback logs that are beyond the retention period. This avoids the overhead of recreating the flashback logs, if another peak workload occurs soon after. Correct answer: A,C
upvoted 21 times

nyanyanyao0826 Highly Voted 6 years, 1 month ago
Correct answer A,C
upvoted 8 times

auwia Most Recent 2 years, 1 month ago
Selected Answer: AD
A and D are true. Oracle until the end will try to don't delete flashback logs. "Proactively" makes sentences false. D is absolutely possible, you can selected on data dictionary tables and set up an event trigger (ex. Database logon, DB scheduler, etc.) to purge logs, but it is not recommended.
upvoted 1 times

auwia 2 years, 2 months ago
Selected Answer: AB
A and B for me.

The oldest flashback logs may be deleted if there is space pressure in the FRA where the flashback logs are stored.

https://docs.oracle.com/en/database/oracle/oracle-database/23/haovw/oracle-flashback-best-practices.html#GUID-3E255563-B001-4AC1-A2E0-1884BF58D6C0
upvoted 1 times

ScottL 3 years, 3 months ago
Selected Answer: AC
A C correct
upvoted 1 times

gio 3 years, 4 months ago
Selected Answer: AC

AC https://docs.oracle.com/en/database/oracle/oracle-database/19/bradv/maintaining-rman-backups.html#GUID-26C1A1E8-5555-4399-9DD1-39BA90E45A0C

upvoted 1 times

mikom80 3 years, 6 months ago

Selected Answer: AC

19c (New Feature) : Auto Space Management for Flashback Logs in the Fast Recovery Area (Doc ID 2465562.1)

upvoted 2 times

rashedptdba 4 years, 10 months ago

Correc Ans:A,C as per 19c (New Feature) : Auto Space Management for Flashback Logs in the Fast Recovery Area (Doc ID 2465562.1)

upvoted 1 times

Dhanushbh 5 years, 1 month ago

agreed with AC

upvoted 2 times

RinD 5 years, 4 months ago

AC

19c (New Feature) : Auto Space Management for Flashback Logs in the Fast Recovery Area (Doc ID 2465562.1)

upvoted 3 times

Jatindra 5 years, 5 months ago

correct ans: AC Until Oracle Database 19c, flashback logs that are beyond the flashback retention period set through the db_flashback_retention_target initialization parameter are marked as reclaimable but are not purged until there is space pressure. The reclaimable flashback logs are deleted only when the fast recovery area is running out of space. The DBA cannot manually delete the flashback logs. This means that when a backup needs to be completed and the FRA is 100 % full, the backup operation must wait until the reclaimable flashback logs are purged to free space in the FRA. With Oracle Database 19c, flashback logs beyond the retention period are proactively deleted without degrading the flashback performance and before there is space pressure. The flashback log space is tracked and monitored proactively to provide flashback log retention in a transparent manner.

upvoted 3 times

auwia 2 years, 2 months ago

So they are A, B corrects and not C!? C is saying BEFORE the space pressure, B is saying AFTER instead.

upvoted 1 times

Chansi 5 years, 11 months ago

I think its A and C. starting from 19c. Starting with Oracle Database Release 19c, Oracle Database monitors flashback logs in the fast recovery area and automatically deletes flashback logs that are beyond the retention period (Not Waiting for FRA Crunch). When the retention target is reduced, flashback logs that are beyond the retention period are deleted immediately.

upvoted 4 times

janw 6 years ago

AB.

https://tekdba.wordpress.com/2019/07/18/oracleflashbackdatabase/#:~:text=The%20Only%20quick%20way%20we,Not%20Waiting%20for%20FRA%20Crunch).

upvoted 1 times

marcinb32 6 years, 1 month ago

correct AB

upvoted 1 times

julica 6 years, 1 month ago

correct A,B

upvoted 2 times

klever35 6 years, 1 month ago

reviewing the indicated by ald85 A, B

upvoted 1 times

ald85 6 years, 2 months ago

I think, A, B https://blogs.oracle.com/maa/flashback-database-and-flashback-logs

upvoted 3 times

Question #27 *Topic 1*

Examine these queries and their output:

After a system crash, an instance restart and an attempted opening of the PDBs result in:

Which two are true? (Choose two.)

```
SQL> select log_mode from v$database;

LOG_MODE
----------------
ARCHIVELOG

SQL> select property_name, property_value
  2 from database_properties where property_name like '%UNDO%';

PROPERTY_NAME            PROPERTY_VALUE
--------------------     --------------------
LOCAL_UNDO_ENABLED FALSE

SQL> select p.name, f.file#, t.name
   2 from v$containers p, v$datafile f, v$tablespace t
   3 where p.con_id=f.con_id
   4 and p.com_id=t.con_id
   5 and t.ts#=f.ts#
   6 order by 1, 2;

NAME        FILE#     NAME
---------  --------  ------------
CDB$ROOT   1         SYSTEM
…
PDB1       24        SYSTEM
…
PDB2       16       SYSTEM
```

```
SQL> startup quiet
ORACLE instance started.
Database mounted.
Database opened.
SQL> alter pluggable database all open;
alter pluggable database all open
*
ERROR at line 1:
ORA-01157: cannot identify/lock data file 24 - see DBWR trace file
ORA-01110: data file 24:
'/u01/oradata/V122CDB1/516000726D464D04E054000C29704164/datafile/o1_mf_system_dmj30kld_.dbf'
```

A. Data file 24 can be recovered while PDB2 is opened.

B. Data file 24 must be recovered while the CDB is opened.

C. Data file 24 can be recovered while CDB$ROOT and PDB$SEED are opened.

D. Data file 24 cannot be recovered while the CDB is opened.

E. Data file 24 must be recovered while PDB2 is closed.

Suggested Answer: *B, D*

Official Answer: BD

Community vote distribution

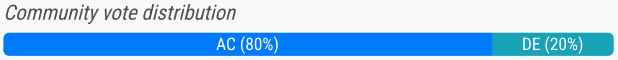

danito Highly Voted 5 years, 10 months ago

D E. Becareful the tablespace afected is system
upvoted 16 times

smartvan Highly Voted 4 years, 9 months ago

A and C.

Datafile 24 belongs to PDB1.

A. It can be recovered while CDB while PDB2 is opened, as it belongs to PDB1.

C. the datafile belongs to PDB1 and CAN be recovered while CDB is opened or in mount state.
upvoted 14 times

auwia Most Recent 2 years, 2 months ago

Selected Answer: DE

Yes: D and E because the tablespace is SYSTEM.
upvoted 1 times

acesonly 2 years, 3 months ago

Selected Answer: AC

TL;DR: I've just tested, the correct answers are A and C. The explanations are below:

A: Since PDB1 is the one missing the data file 24, the recovery is possible while PDB2 is opened, i.e. PDB1 is unrelated to PDB2.

C: The data file 24 can be recovered while the CDB$ROOT and PDB$SEED are opened (or while both are MOUNTED).

(Not B): The data file 24 recovery can also be done while CDB is MOUNTED; option B is too restrictive and, therefore incorrect.

(Not D): For the data file 24 recovery, the CDB should be OPENED or at least MOUNTED, otherwise, the recovery is not possible (when shut down).

(Not E): Since PDB1 is the one missing the data file 24, the recovery is possible regardless of the PDB2; PDB2 may be open or closed, it's all the same, so the option E is too restrictive and, therefore incorrect.
upvoted 2 times

dancymonkey 2 years, 7 months ago

A and C for me. DE > Think simple, Is there necessary to close PDB2 and CDB just to recover PDB1?
upvoted 1 times

Yennism 2 years, 8 months ago

Selected Answer: AC

Datafile 24 belongs to odb PDB1 the datafile belongs to PDB1 and CAN be recovered while CDB is opened or in mount
upvoted 1 times

Emperor11 2 years, 8 months ago

Selected Answer: AC

ERROR en linea 1: ORA-01157: no se puede identificar/bloquear el archivo de datos 70 - consulte el archivo de rastreo del DBWR ORA-01110: archivo de datos 70: '/u02/app/oracle/oradata/ORCL/pdb1/ORCL/0B35DE71576F30A2E063F038A8C0 6E99/datafile/o1_mf_system_lpcnc7r6_.dbf'

SQL> select name, open_mode from v$containers;

NAME OPEN_MODE ---------- ---------- CDB$ROOT READ WRITE PDB$SEED READ ONLY NEWPDB READ WRITE PDB1 MOUNTED

restore datafile 70; recover datafile 70; SQL> alter pluggable database all open;

Base de datos de conexion modificada.

SQL> show pdbs

CON_ID CON_NAME OPEN MODE RESTRICTED --------- ---------- --------- --------- 2 PDB$SEED READ ONLY NO 3 NEWPDB READ WRITE NO 5 PDB1 READ WRITE NO
upvoted 2 times

LEOC71 2 years, 8 months ago

Selected Answer: AC

Tested: 19c RU 19.20 SQL > startup SQL> alter pluggable database all open; ERROR at line 1: ORA-01157: cannot identify/lock data file 9 - see DBWR trace file ORA-01110: data file 9: '/u02/oradata/CDB2/PDB10/system01.dbf' RMAN > restore datafile 9; RMAN> recover datafile 9; SQL> ALTER PLUGGABLE DATABASE pdb10 OPEN;

upvoted 1 times

musafir 2 years, 11 months ago

Selected Answer: DE

D E Since it's SYSTEM tablespace, otherwise it would be A C

upvoted 1 times

gio 3 years, 4 months ago

Selected Answer: DE

SYSTEM tablespace so CDB must be in mount state

upvoted 1 times

jareach 3 years, 6 months ago

Selected Answer: AC

Suggested BD are mutually exclusive.

upvoted 1 times

Manojx 3 years, 7 months ago

A and C are correct. It doesn't matter if tablespace to be recovered in SYSTEM. https://docs.oracle.com/en/database/oracle/oracle-database/19/bradv/rman-performing-flashback-dbpitr.html#GUID-35D290AC-CC01-4984-93D6-AE41A3D0D3B6

upvoted 1 times

hilaire 3 years, 11 months ago

AC is the correct answer i test it

upvoted 1 times

Halewar 4 years, 3 months ago

Selected Answer: AC

AC for me

upvoted 5 times

Yennism 4 years, 7 months ago

A, C. Remember that datafile belongs to PDB1. And CDB can be in OPEN or MOUNT state to restore de datafile of PDB2.

upvoted 7 times

chefdba 4 years, 9 months ago

AC = correct answers Take note: DATA FILE 24 is for PB1

upvoted 3 times

rashedptdba 4 years, 10 months ago

Ans: is A,C. Tested on my LAB.

upvoted 6 times

Question #28 *Topic 1*

Which two are true about RMAN duplexed backup sets? (Choose two.)

A. A duplexed backup set uses the same number of SBT channels as a non-duplexed backup set for the same number of files.

B. A non-duplexed backup set written to disk can be duplexed to disk by backing up the backup set that is already on disk.

C. A non-duplexed backup set written to SBT can be duplexed to tape by backing up the backup set that is already on tape.

D. A non-duplexed backup set written to disk can be duplexed to tape by backing up the backup set that is already on disk.

E. A non-duplexed backup set written to SBT can be duplexed to disk by backing up the backup set that is already on tape.

F. A duplexed backup set always uses twice as many SBT channels as a non-duplexed backup set for the same number of files.

Suggested Answer: *D, F*

Official Answer: DF

Community vote distribution



ald85 Highly Voted 5 years, 8 months ago

I think B D -> https://docs.oracle.com/cd/E18283_01/backup.112/e10642/rcmbckad.htm#i1006180

upvoted 17 times

Datajimm Highly Voted 3 years, 10 months ago

Selected Answer: BD

"You cannot back up from tape to tape or from tape to disk: only from disk to disk or disk to tape."

https://docs.oracle.com/en/database/oracle/oracle-database/19/rcmrf/BACKUP.html#GUID-73642FF2-43C5-48B2-9969-99001C52EB50

upvoted 5 times

dancymonkey Most Recent 2 years ago

Agreed BD

upvoted 1 times

Guhborges 2 years, 6 months ago

Selected Answer: BD

B D -> You cannot back up from tape to tape or from tape to disk

upvoted 2 times

12345aws 2 years, 6 months ago

B & D are correct

upvoted 1 times

kaka321 2 years, 7 months ago

disk to disk, tape to tape BC

upvoted 1 times

gio 2 years, 10 months ago

Selected Answer: BC

BACKUP COPIES command work tape to tape or disk to disk

upvoted 1 times

jareach 3 years ago

Selected Answer: BC

Backups cannot be duplexed on tape/disk combination.

upvoted 1 times

Marco_M 4 years, 9 months ago

I think B D

https://web.stanford.edu/dept/itss/docs/oracle/10gR2/backup.102/b14191/rcmconc1003.htm#:~:text=Backups%20of%20Backup%20Sets&text=Only%20backup%20sets%20that%20were,run%20BACKUP%20AS%20COPY%20BACKUPSET%20.

Backups of Backup Sets The RMAN BACKUP BACKUPSET command backs up previously created backup sets. Only backup sets that were created on device type DISK can be backed up using RMAN. They can be backed up to any available device type.

upvoted 4 times

RinD 4 years, 11 months ago

B, D

https://docs.oracle.com/cd/E18283_01/backup.112/e10642/rcmcncpt.htm#:~:text=The%20RMAN%20BACKUP%20BACKUPSET%20command,of%20the%20same%20backup%20set.

upvoted 4 times

ama 5 years ago

I think B, C are correct, at least according to the provided explaination from Oracle:

If your local backup uses a disk channel and want to Duplex to Cloud, Cloud backup uses and SBT channel. The same backup can not be duplexed to mixed kind of channels, see below. Duplexing Backup Sets RMAN can duplex backups to either disk or tape, but cannot duplex backups to tape and disk simultaneously.

upvoted 4 times

pgtips 5 years, 5 months ago

The RMAN BACKUP BACKUPSET command allows you to do both, 1) Disk backup set to TAPE, 2) Disk Backup set to DISK. B. non-duplexed backup set written to disk can be duplexed to disk by backing up the backup set that is already on disk. D. A non-duplexed backup set written to disk can be duplexed to tape by backing up the backup set that is already on disk

upvoted 1 times

Jatindra 5 years, 8 months ago

F option is not there in actual exam.

upvoted 2 times

veginha 5 years, 8 months ago

D seems to be right, which is the other one that is right?

upvoted 1 times

Question #29 *Topic 1*

Which three are true about RMAN persistent configuration settings, administration, and their effects? (Choose three.)

A. A target database's persistent RMAN configuration settings are always stored in the target's control file

B. Backup older than the recovery window retention policy are always deleted automatically if the backup location has insufficient space.

C. Backups written to the fast recovery area (FRA) that are obsolete based on the redundancy retention policy can be deleted automatically to free space.

D. The RMAN SHOW ALL command displays only settings with nondefault values.

E. A target database's persistent RMAN configuration settings are always synchronized automatically with the RMAN catalog.

F. The V$RMAN_CONFIGURATION view displays only settings with values that have been modified.

G. A DBA must specify either a redundancy retention policy or a recovery window retention policy.

Suggested Answer: *E, F, G*
Official Answer: EFG

Community vote distribution
ACF (100%)

RinD Highly Voted 3 years, 11 months ago
Agree with ACF
upvoted 17 times

michael_live Highly Voted 4 years, 4 months ago
Correct answer is ACF - tested all in lab.
upvoted 7 times

ScottL Most Recent 1 year, 9 months ago
Selected Answer: ACF
ACF correct
upvoted 1 times

gio 1 year, 9 months ago
Selected Answer: ACF
F is correct, I try to modify an RMAN configuration and v$rman_configuration shows only the modification. A is correct because information are always stored in control file even if you use rman catalog C is correct (delete obsolete)
upvoted 2 times

jareach 2 years ago
Selected Answer: ACF
Retention policy -- as suggested -- does not have to be specified, whereas the persistent settings are always stored in the target db controlfile.
upvoted 2 times

hilaire 2 years, 5 months ago
ACF correct
upvoted 2 times

Yennism 3 years, 1 month ago
Agree with ACF
upvoted 3 times

kameni 3 years, 5 months ago
I would say correct is A,C,E E is correct --> RMAN performs resynchronizations automatically as needed when you execute certain commands, including BACKUP. You can also manually perform a full resynchronization using the RESYNC CATALOG command. F is incorrect --> V$RMAN_CONFIGURATION lists information about RMAN persistent configuration settings.
upvoted 3 times

wqui 3 years, 5 months ago
As you said, RMAN performs resync config in some comands, not always
upvoted 1 times

PortlandFighters 3 years, 4 months ago

Agree F should be incorrect. As stated in Oracle 19c database reference, V$RMAN_CONFIGURATION lists information about RMAN persistent configuration settings. https://docs.oracle.com/en/database/oracle/oracle-database/19/refrn/V-RMAN_CONFIGURATION.html#GUID-13D7B8D1-C7B7-43EC-B176-90CCB9A4D647

upvoted 1 times

Kishor123 3 years, 6 months ago

AFG C: RMAN backup in FRA never get deleted AUTOMATICALLY. they have to either delete manually or using script based on retention policy.Whereas Flashback logs get automatically deleted based on availability of free space or if FRA parameter db_flashback_retention_target value reduced(19c feature).

upvoted 1 times

brolem 2 years, 10 months ago

C is correct : If you use a fast recovery area as your only disk-based backup destination, then you never have to delete obsolete backups from disk. The fast recovery area keeps files as specified by the retention policy, and deletes them only when space is needed.

upvoted 3 times

xRodge 3 years, 6 months ago

G cannot be correct - as Oracle will provide a default value, when a DB is created. The DBA does not "need" to specify either of these values.

I think that C is not written perfectly, but it does imply that the use of the "DELETE OBSOLETE" statement will automatically remove every non-needed backup file that lies outside of the retention policy period.

So, I believe that ACF is correct.

upvoted 1 times

Kishor123 3 years, 6 months ago

Got it. Thanks ! Agree with ACF

upvoted 1 times

pgtips 4 years, 6 months ago

ACF sounds more relevant.

upvoted 4 times

ald85 4 years, 8 months ago

I think CFG

upvoted 2 times

veginha 4 years, 8 months ago

ACF Why does the dba need to specify the retention policy? it's not necessary, you can use the default values also.

upvoted 7 times

Question #30 *Topic 1*

Which three are true about Optimizer Statistics Advisor? (Choose three.)

A. It can be run only manually.

B. It is part of the DBMS_ADVISOR package.

C. It can recommend changes to improve the statistics gathering process.

D. It always analyzes all schemas in the database.

E. It runs automatically every night by default.

F. It is part of the DBMS_STATS package.

Suggested Answer: *C, D, F*

Reference: https://oracle-base.com/articles/12c/optimizer-statistics-advisor-12cr2

Community vote distribution

CEF (100%)

monad2006 Highly Voted 5 years, 2 months ago

I think: CEF

upvoted 19 times

Guhborges Most Recent 2 years, 1 month ago

Selected Answer: CEF

C,E,F.

The advisor task runs automatically in the maintenance window, but you can also run it on demand.

https://docs.oracle.com/en/database/oracle/oracle-database/21/tgsql/optimizer-statistics-advisor.html#GUID-11B3764D-CCA1-415F-88DD-0F6884DF8D89

upvoted 2 times

ScottL 2 years, 3 months ago

Selected Answer: CEF

CEF correct

upvoted 1 times

gio 2 years, 3 months ago

Selected Answer: CEF

I think cef

upvoted 1 times

vkra 2 years, 6 months ago

From: https://docs.oracle.com/en/database/oracle/oracle-database/21/arpls/DBMS_STATS.html#GUID-A405E9B8-B8C8-400D-A838-F7DA161BE1A2
Optimizer Statistics Advisor inspects the statistics gathering process, automatically diagnoses problems in the existing practices for gathering statistics, and then generates a report of findings and recommendations. The advisor task runs automatically in the maintenance window. However, you can also run the job on demand.

upvoted 1 times

vkra 2 years, 6 months ago

Selected Answer: CEF

agree with CEF

upvoted 1 times

Dhanushbh 4 years, 1 month ago

agree with CEF

upvoted 4 times

Heat418 3 years, 1 month ago

I would go with CEF

upvoted 2 times

Adel_007 4 years, 10 months ago

Optimizer Statistics Advisor is built-in diagnostic software that analyzes the quality of statistics and statistics-related tasks. The advisor task runs automatically in the maintenance window, but you can also run it on demand. You can then view the advisor report. If the advisor makes recommendations, then in some cases you can run system-generated scripts to implement them. his section contains the following topics: Optimizer Statistics Advisor inspects how optimizer statistics are gathered. Optimizer Statistics Advisor uses the same advisor framework as Automatic Database Diagnostic Monitor (ADDM), SQL Performance Analyzer, and other advisors. Perform Optimizer Statistics Advisor tasks using the DBMS_STATS PL/SQL package. https://docs.oracle.com/en/database/oracle/oracle-database/12.2/tgsql/optimizer-statistics-advisor.html#GUID-D81A7708-FDA0-45BB-A6E2-103858B047AE

upvoted 1 times

pgtips 5 years ago

CEF - SOUNDS CORRECT

upvoted 2 times

cerebro2000x 5 years ago

AUTO_STATS_ADVISOR_TASK

The Optimizer Statistics Advisor runs a task called AUTO_STATS_ADVISOR_TASK during the maintenance window, so by default you don't need to do anything to start seeing advice about your database statistics.

upvoted 1 times

cerebro2000x 5 years ago

DECLARE tname VARCHAR2(32767) := 'demo'; -- task name BEGIN tname := dbms_stats.create_advisor_task(tname); END; / DECLARE tname VARCHAR2(32767) := 'demo'; -- task name ename VARCHAR2(32767) := NULL; -- execute name BEGIN ename := dbms_stats.execute_advisor_task(tname); END; /

upvoted 2 times

dacoben415lywenw 5 years, 1 month ago

CEF, is the only right answer. A is not correct, daily maintenance window starts at 10pm and spans the whole night.

upvoted 2 times

julica 5 years, 1 month ago

CEF The advisor task ('Statistics Advisor') runs automatically in the maintenance window. And of course you can also run it on demand

upvoted 2 times

Jatindra 5 years, 2 months ago

Correct Answer: ACF https://mikedietrichde.com/2017/08/22/oracle-optimizer-statistics-advisor-in-oracle-database-12-2-0-1/ https://www.oracle.com/technetwork/database/bi-datawarehousing/twp-bp-for-stats-gather-19c-5324205.pdf

upvoted 3 times

ObserverPL 4 years, 3 months ago

Mate, you provided link to Mike's post and there is clearly The advisor task ('Statistics Advisor') runs automatically in the maintenance window. so A is not correct, come on...

upvoted 3 times

navingupta52 4 years, 1 month ago

He always posts wrong answers and links with confidence.

upvoted 3 times

Question #31 *Topic 1*

You issued this command:
RMAN> BACKUP RECOVERY FILES;
Which two are true? (Choose two.)

A. All Oracle recovery files not in the current FRA that have not been backed up already, are backed up.

B. All non-Oracle files in the current FRA that have not been backed up already, are backed up.

C. All Oracle recovery files in the current FRA that have not been backed up already, are backed up.

D. All Oracle recovery files in the current fast recovery area (FRA) are backed up.

E. These backups can be written to disk or SBT.

Suggested Answer: *B, E*
Official Answer: BE

Community vote distribution
AC (78%) DE (22%)

monad2006 Highly Voted 5 years, 8 months ago
I think: AC
upvoted 20 times

proact 5 years, 4 months ago
A is wrong as USER_TABLESPACE doesn't specify a default tablespace, but specifies which tablespaces are available in the new PDB.
upvoted 1 times

SimoneF Highly Voted 4 years, 10 months ago
https://docs.oracle.com/en/database/oracle/oracle-database/19/sbydb/using-RMAN-in-oracle-data-guard-configurations.html#GUID-66665A86-0B71-41FE-BED0-B9992A405C35

"To back up all recovery-related files to tape, use the RMAN BACKUP RECOVERY FILES [...] This ensures that all current incremental, image copy, and archived log backups on disk are backed up to tape."

and https://docs.oracle.com/en/database/oracle/oracle-database/19/bradv/configuring-rman-client-basic.html#GUID-ABA3E439-57C2-4AE7-BCEF-39A6663E7AA2

"RMAN uses backup optimization when the following conditions are true: [...] You run [...] BACKUP RECOVERY FILES [...]"

Lead me to believe AC
upvoted 7 times

Heat418 3 years, 7 months ago
I vote for AC as well.
upvoted 2 times

SimoneF 4 years, 10 months ago
Sorry for the triple post, it didn't give me immediate results so I thought publication failed.

Anyway, in the first case, it's said that ALL of those files on disk are backed up to tape, so also outside the recovery area (as opposed to "BACKUP RECOVERY AREA").

The second sentence states that backup optimization is turned on with BACKUP RECOVERY FILES: this implies that "the BACKUP command skips backing up files when the identical file has been backed up to the specified device type."

Since E is wrong cause you need to specify "TO LOCATION" in order to back up to disk in this case:

A and C should be right.

upvoted 2 times

1451755 Most Recent 1 year, 4 months ago

Selected Answer: AC

ee https://docs.oracle.com/en/database/oracle/oracle-database/19/rcmrf/BACKUP.html#GUID-73642FF2-43C5-48B2-9969-99001C52EB50 "BACKUP RECOVERY FILES Backs up all recovery files on disk, whether they are stored in the fast recovery area or other locations on disk. The backups can go to SBT or disk. To backup to disk, you must use the TO DESTINATION syntax ... Recovery files include full and incremental backup sets, control file autobackups, archived redo log files, and data file copies. By default, backup optimization is enabled for this command even if the CONFIGURE BACKUP OPTIMIZATION setting is OFF.

Conclusion: A & C true: Backup whether they are stored in the fast recovery area or other locations on disk (backup optimization is enabled = have not been backed up already). B false: non-Oracle files are not recovery files D false: FORCE syntax not present E false: TO DESTINATION syntax not present, this command is for SBT cannot be written to disk with this command

upvoted 1 times

antonica 1 year, 11 months ago

an example of recovery files? please

upvoted 1 times

Shunpin 2 years, 2 months ago

Selected Answer: AC

https://docs.oracle.com/en/database/oracle/oracle-database/19/rcmrf/BACKUP.html#GUID-73642FF2-43C5-48B2-9969-99001C52EB50

Backs up all recovery files on disk, whether they are stored in the fast recovery area or other locations on disk The backups can go to SBT or disk E might have questions - If the backups are not on disk.

upvoted 1 times

LEOC71 2 years, 2 months ago

Selected Answer: AC

"Backs up all recovery files on disk, whether they are stored in the fast recovery area or other locations on disk. " "By default, backup optimization is enabled for this command even if the CONFIGURE BACKUP OPTIMIZATION setting is OFF" => A:TRUE, C:TRUE

upvoted 1 times

piontk 2 years, 6 months ago

Selected Answer: AC

A, C are correct:

About BACKUP RECOVERY FILES: "Backs up all recovery files on disk, whether they are stored in the fast recovery area or other locations on disk. [...] By default, backup optimization is enabled for this command even if the CONFIGURE BACKUP OPTIMIZATION setting is OFF. "

About BACKUP OPTIMIZATION: "Optimization prevents RMAN from backing up a file to a device type if the identical file is already backed up on the device type".

The point that invalidates E is: "The backups can go to SBT or disk. To backup to disk, you must use the TO DESTINATION syntax outlined in toDestSpec. " Since in this question we do not use TO DESTINATION, we are not backing up to disk.

REF: https://docs.oracle.com/en/database/oracle/oracle-database/19/rcmrf/BACKUP.html#d11234e4347
https://docs.oracle.com/en/database/oracle/oracle-database/19/rcmrf/CONFIGURE.html#d27691e2023

upvoted 1 times

kaka321 2 years, 7 months ago

I will go for DE. E is 100% correct backup recovery files; can be done on any of tape or disk. its 100% right that it backs up all oracle files in FRA in the document it does not specified leaving already backed up files. https://docs.oracle.com/en/database/oracle/oracle-database/19/rcmrf/BACKUP.html#GUID-73642FF2-43C5-48B2-9969-99001C52EB50

upvoted 1 times

ErikJanssen 2 years, 6 months ago

E is also wrong I think; When backing up to disk, one needs to specify a toDestination clause which is not mentioned.

upvoted 1 times

ErikJanssen 2 years, 6 months ago

I think D is wrong; The command enables Rman backup optimization; So only the files not already backed u will be backed up and not all files. See: https://docs.oracle.com/en/database/oracle/oracle-database/19/rcmrf/BACKUP.html#GUID-73642FF2-43C5-48B2-9969-99001C52EB50

upvoted 1 times

Guhborges 2 years, 7 months ago

Selected Answer: DE

D, E 100%

upvoted 1 times

ScottL 2 years, 9 months ago

Selected Answer: AC

A,C Should be correct

upvoted 1 times

gio 2 years, 9 months ago

Selected Answer: AC

" Backs up all recovery files on disk, whether they are stored in the fast recovery area or other locations on disk. The backups can go to SBT or disk. To backup to disk, you must use the TO DESTINATION syntax outlined in toDestSpec.

Recovery files include full and incremental backup sets, control file autobackups, archived redo log files, and data file copies.

By default, backup optimization is enabled for this command even if the CONFIGURE BACKUP OPTIMIZATION setting is OFF. You can disable backup optimization for RECOVERY FILES by specifying FORCE. " from: https://docs.oracle.com/en/database/oracle/oracle-database/19/rcmrf/BACKUP.html#GUID-73642FF2-43C5-48B2-9969-99001C52EB50

upvoted 1 times

ErikJanssen 2 years, 10 months ago

I found this piece of documentation about backing up RMAN backups very clarifying: https://docs.oracle.com/en/database/oracle/oracle-database/19/bradv/backing-up-database.html#GUID-731ADA67-235E-481D-AC17-9F9067AD55BF

upvoted 1 times

vkra 3 years ago

Selected Answer: AC

AC should be correct

upvoted 1 times

hilaire 3 years, 1 month ago

AC correct

upvoted 1 times

hilaire 3 years, 5 months ago

AC coorect

upvoted 1 times

flaviogcmelo 3 years, 10 months ago

Selected Answer: DE

In my opinion files already backed up will be disregarded by the optimized backup function because FORCE was not used. And the backup can be sent either to disk or tape, in this case it will be to tape, if there was the TO DESTINATION clause it would be written to disk. Reference: https://docs.oracle.com/en/database/oracle/oracle-database/19/rcmrf/BACKUP.html#GUID-73642FF2-43C5-48B2-9969-99001C52EB50 Backs up all recovery files on disk, whether they are stored in the fast recovery area or other locations on disk. The backups can go to SBT or disk. To backup to disk, you must use the TO DESTINATION syntax outlined in toDestSpec. By default, backup optimization is enabled for this command even if the CONFIGURE BACKUP OPTIMIZATION setting is OFF. You can disable backup optimization for RECOVERY FILES by specifying FORCE. In my understanding this mean that only files not backed up will be added.

upvoted 1 times

MartinDBA2 3 years, 10 months ago

D is the only correct answer, I have a problem with A&C which states "not backed up yet" as the manual says "all recovery files" which indicates even if they have been backup up previously.

https://docs.oracle.com/en/database/oracle/oracle-database/19/rcmrf/BACKUP.html#GUID-73642FF2-43C5-48B2-9969-99001C52EB50 Backs up all recovery files on disk, whether they are stored in the fast recovery area or other locations on disk. The backups can go to SBT or disk. To backup to disk, you must use the TO DESTINATION syntax outlined in toDestSpec.

E is also incorrect as the command given does not include the necessary TO DESTINATION syntax

upvoted 2 times

FelipeC 2 years, 6 months ago

if you I have a problem with A&C which states "not backed up yet", think about D which states "ALL" ...., flashback logs are never backed up

upvoted 1 times

Question #32 *Topic 1*

Which two are true about the Oracle database methodology? (Choose two.)

A. The Oracle Database time model should be used to find the database and instance areas most in need of tuning.

B. Tuning activities should stop once the user is satisfied with performance.

C. Tuning activities should stop once agreed service levels for performance have been met.

D. The database instance memory should always be tuned before tuning any file systems.

E. SQL statements should always be tuned before tuning any file systems.

F. The alert log should be used to find the database and instance areas most in need of tuning.

Suggested Answer: *C, F*
Reference: https://flylib.com/books/en/4.322.1.9/1/

janw Highly Voted 6 years ago
AE i suppose
upvoted 13 times

xRodge Highly Voted 5 years ago
Few differences of opinions here.

I believe that A is definitely correct. It will show what the database is spending most time on, that (usually) indicates where tuning is required. B and C relate to user experience. We (DBAs) should work with users, to help achieve their requirements. However, we can be proactive outside of this. Just because a user is happy - doesn't mean we can't make them "happier" by doing tuning activities before things get bad. So, neither of these should be true. D is incorrect, memory is much less likely to be a culprit of an issue, than a filesystem I/O issue. E is correct. We should never try to tune a file system first. If we make a file system change that fixes one SQL statement - it might harm a different SQL statement. However, if we tune the first SQL statement, it will have no impact on the rest of the system. Therefore, we should try to tune "bad SQL" first. F is incorrect - the alert log gives us little/no information to do with tuning.

So, AE are correct.
upvoted 10 times

acesonly 2 years, 3 months ago
Great explanation, thank you. I personally agree, however, I'm under impression that in some cases, correct answers don't necessarily align with the common sense, but are based in an obscure sentence in the documentation, as here: "The Oracle performance method can be applied until performance goals are met or deemed impractical." https://docs.oracle.com/en/database/oracle/oracle-database/19/tdppt/oracle-database-performance-method.html#GUID-811C438C-149B-4BFE-A40A-1FAF2BF67C95:~:text=The%20Oracle%20performance%20method%20can%20be%20applied%20until%20performance%20goals%20are%20met%20or%20deemed%20impractical.
upvoted 1 times

acesonly 2 years, 3 months ago
I doubt that some answers are not offered in good faith, such as option C: "Tuning activities should stop once agreed service levels for performance have been met.". I'm aiming at the "...should..." vs. the "...can..." in the link I gave above.
upvoted 1 times

minkpang Most Recent 1 year ago
Selected Answer: AE
A,E is correct
upvoted 1 times

1451755 1 year, 4 months ago
Selected Answer: AE
We can/shlud tune above SLA, if it is make sence.
upvoted 2 times

wagihov 1 year, 8 months ago
Selected Answer: AE

Following oracle tuning methodology for tuning from top to bottom, E is correct and D is wrong. B and C are almost the same and very tricky. Although oracle methodology is to stop tuning when a goal is reached, but this goal shouldn't be just a SLA or a satisfied user as the goal can be beyond that. A is definitely correct, while F is wrong as alert.log is irrelevant to performance tuning.

upvoted 1 times

cujar2003 1 year, 9 months ago

Selected Answer: AE

The Oracle Database time model should be used to find the database and instance areas most in need of tuning. SQL statements should always be tuned before tuning any file systems.

upvoted 1 times

hojunlee 2 years, 10 months ago

I guess that correct answers are A and C.

upvoted 1 times

flaviogcmelo 2 years, 11 months ago

Selected Answer: AC

A. https://docs.oracle.com/en/database/oracle/oracle-database/19/tdppt/oracle-database-performance-method.html#GUID-8DEA5079-6AAB-47D8-857A-9B5301843290 C. https://docs.oracle.com/en/database/oracle/oracle-database/19/tdppt/oracle-database-performance-method.html#GUID-6203E485-FA06-4EE3-9003-4449FBDE214B

upvoted 1 times

kaka321 3 years, 1 month ago

AC is right tunning should be based on agreed SLA see https://docs.oracle.com/en/database/oracle/oracle-database/19/tdppt/oracle-database-performance-method.html#GUID-D5B54332-FE40-4537-B1C9-0704343CEF54

upvoted 1 times

Guhborges 3 years, 1 month ago

Selected Answer: AC

The two true statements about the Oracle database methodology are:

A. The Oracle Database time model should be used to find the database and instance areas most in need of tuning. C. Tuning activities should stop once agreed service levels for performance have been met.

Why not E?

E. While tuning SQL statements is an important aspect of Oracle database performance tuning, it is not necessary to tune SQL statements before tuning any file systems. Again, the order of tuning activities may depend on the specific needs of the system.

upvoted 1 times

Guhborges 3 years, 1 month ago

Selected Answer: AC

A,C 100% sure

upvoted 1 times

gio 3 years, 3 months ago

Selected Answer: AC

https://docs.oracle.com/en/database/oracle/oracle-database/19/tdppt/oracle-database-performance-method.html#GUID-6203E485-FA06-4EE3-9003-4449FBDE214B

upvoted 1 times

jareach 3 years, 6 months ago

Selected Answer: AE

A: time model gives the best view of how the instance spends its time. E: Unless the instance configuration is totally off, the SQL statement tuning has the greatest performance benefit.

upvoted 2 times

G_C 3 years, 10 months ago

Selected Answer: AC

Because of the same reason as Asmodeus: 1. This feedback makes it easier to set performance goals. Improved performance can be measured in terms of business goals rather than system statistics (...) 2. (...) The Oracle performance method can be applied until performance goals are met or deemed impractical (...).

upvoted 1 times

Shidt 4 years, 2 months ago

NIEMNkaoguolema

upvoted 1 times

brolem 4 years, 3 months ago

The alert log and trace files are vital in instance tuning. They may point you to the problem (memory bottlenecks, locked sessions): - End to End Application tracing identifies excessive workloads on the system by specific user, service, or application component. - Server-generated alerts automatically provide notifications when impending problems are detected. See Oracle Database Administrator's Guide to learn how to monitor the operation of the database with server-generated alerts.

upvoted 1 times

Neil107 5 years ago

It's AE. Tuning is an ongoing process.

upvoted 2 times

Question #33 *Topic 1*

While backing up to the Oracle Fast Recovery Area (FRA), you determined the backup is taking too long and suspect a performance bottleneck. Which three are true about diagnosing and tuning these problems? (Choose three.)

A. If an RMAN BACKUP VALIDATE command takes roughly the same time as an actual backup, then both read and write I/O are likely bottlenecks.

B. Setting DBWR_IO_SLAVES to a non zero value can improve backup performance when using synchronous I/O.

C. If an RMAN BACKUP VALIDATE command takes noticeably less than an actual backup, then write I/O is a likely bottleneck.

D. If an RMAN BACKUP VALIDATE command takes roughly the same time as an actual backup, then read I/O is a likely bottleneck.

E. Data files with a high value in V$BACKUP_SYNC_IO.DISCRETE_BYTES_PER_SECOND are a potential performance bottleneck when synchronous I/O is used.

F. Setting DBWR_IO_SLAVES to a non zero value can improve backup performance when using asynchronous I/O/

G. Data files with a high value in V$BACKUP_ASYNC_IO.SHORT_WAITS are a potential performance bottleneck when asynchronous I/O is used.

Suggested Answer: *B, C, E*

Reference: https://web.stanford.edu/dept/itss/docs/oracle/10gR2/backup.102/b14191/rcmtunin003.htm

Community vote distribution

BCD (100%)

ssseee Highly Voted 3 years, 9 months ago

http://docshare04.docshare.tips/files/21043/210437276.pdf

correct answer BCD.

check above document

upvoted 13 times

kameni Highly Voted 3 years, 5 months ago

Answer: B,C,D

One reliable way to determine whether the output device or input disk I/O is the bottleneck in a given backup job is to compare the time required to run backup tasks with the time required to run BACKUP VALIDATE of the same tasks. BACKUP VALIDATE of a backup performs the same disk reads as a real backup but performs no I/O to an output device. D -> If the time for the BACKUP VALIDATE to tape is about the same as the time for a real backup to tape, then reading from disk is the likely bottleneck. C -> If the time for the BACKUP VALIDATE to tape is significantly less than the time for a real backup to tape, then writing to the output device is the likely bottleneck. B -> Some operating systems support native asynchronous I/O. If and only if your disk does not support asynchronous I/O, then set DBWR_IO_SLAVES. Any nonzero value for DBWR_IO_SLAVES causes a fixed number of disk I/O slaves to be used for backup and restore, which simulates asynchronous I/O.

G -> incorrect - The simplest way to identify the bottleneck is to find the data file that has the largest ratio for LONG_WAITS divided by IO_COUNT

upvoted 8 times

ScottL Most Recent 1 year, 9 months ago

Selected Answer: BCD

BCD Correct

upvoted 1 times

gio 1 year, 9 months ago

Selected Answer: BCD

i think bcd

upvoted 1 times

vkra 2 years ago

Selected Answer: BCD

BCD should be correct

upvoted 1 times

Aldrid 3 years, 8 months ago

BC Correct A,F WRONG D/E/G - in this formulation does not allow to draw an unequivocal conclusion

I go with BCD..

upvoted 2 times

wiprooracle 3 years, 8 months ago

B,C,E

E , as per https://docs.oracle.com/en/database/oracle/oracle-database/18/bradv/tuning-rman-performance.html#GUID-5D8AAB8F-8A63-4B82-9BFB-16D22CE439FD

upvoted 2 times

PortlandFighters 3 years, 4 months ago

Yo mate, your link indicated that E is incorrect instead. The sentence of E is "Data files with a high value in V$BACKUP_SYNC_IO.DISCRETE_BYTES_PER_SECOND are a potential performance bottleneck when synchronous I/O is used". So, a high value means a high throughput with large number of bytes being transferred per second. Besides, the doc ask you to compare the transfer rate with the specified value with the device specification. Say the device support up to 1GB per second but the DISCRETE_BYTES_PER_SECOND is much lower then this value, then there is an issue.

upvoted 2 times

Manojx 3 years, 8 months ago

BCD DBWR_IO_SLAVES is used for asynchronous I/O.

High value of V$BACKUP_SYNC_IO.DISCRETE_BYTES_PER_SECOND is good.
https://docs.oracle.com/cd/B12037_01/server.101/b10755/dynviews_1033.htm

High value in V$BACKUP_ASYNC_IO.SHORT_WAITS is good. Short waits are the number of times the backup or restore process made an operating system call to poll for I/O completion in a nonblocking mode. https://docs.oracle.com/cd/E18283_01/backup.112/e10642/rcmtunin.htm

upvoted 1 times

asefa 4 years ago

Student GUide SET DBWR_IO_SLAVES if you use synchronous disk I/. if your disk does not support asnychronous I/O, try setting DBWR_IO_SLAVES initilization parameter to a nonzero value. any nonzero value for DBWR_IO_SLAVES causes fized number of 4 of disk I/O slaves to be used for backup and restore, simulating ASYNC IO if backup validate time is less than the actual backup time, buffer copy or write to storage is the likely bottlneck if BACKUP validate time ~= actual backup time, the read phase is the likely bottleneck. so what do u think but ill go with B,C,D

upvoted 3 times

Schmeili 4 years, 3 months ago

E is wrong:

Determining Bottlenecks with Synchronous I/O With synchronous I/O, it is difficult to identify specific bottlenecks because all synchronous I/O is a bottleneck to the process. The only way to tune synchronous I/O is to compare the bytes-per-second rate with the device's maximum throughput rate. If the bytes-per-second rate is lower than that device specifies, consider tuning that part of the backup/restore process. Use the V$BACKUP_SYNC_IO.DISCRETE_BYTES_PER_SECOND column to see the I/O rate. Source:
https://www.csee.umbc.edu/portal/help/oracle8/server.815/a67775/ch20_io.htm

upvoted 1 times

undoundo 4 years, 3 months ago

Why is D wrong?

upvoted 2 times

bugzbinny 4 years, 4 months ago

B Set DBWR_IO_SLAVES if you use synchronous disk I/O. If your disk does not support asynchronous I/O, then try setting the DBWR_IO_SLAVES initialization parameter to a nonzero value C Issue a backup validate command to provide the time required to read data files. Compare the time that command takes to execute to the time it takes to perform a regular backup. The difference will give you the time RMAN is spending on writing. E Query the DISCRETE_BYTES_PER_SECOND column from V$BACKUP_SYNC_IO to view the I/O rate

upvoted 3 times

proact 4 years, 4 months ago

Corrects are BCE. It is B correct as: If you set DBWR_IO_SLAVES to a nonzero value,..... the number of I/O server processes used by Recovery Manager is set to 4 only if asynchronous I/O is disabled (either your platform does not support asynchronous I/O or disk_asynch_io is set to false).

https://docs.oracle.com/en/database/oracle/oracle-database/19/refrn/DBWR_IO_SLAVES.html#GUID-D8A04F0F-0D5A-4B28-895E-8F63BD3B7DC9

So, when DBWR_IO_SLAVES is set to a nonzero value, RMAN is set to 4 only if configuration is synchronous I/O.

upvoted 1 times

Chansi 4 years, 5 months ago

I think its bcd http://docshare04.docshare.tips/files/21043/210437276.pdf search for bottleneck

upvoted 1 times

janw 4 years, 6 months ago

If the time for the BACKUP VALIDATE to tape is about the same as the time for a real backup to tape, then reading from disk is the likely bottleneck. See "Tuning the Read Phase".

If the time for the BACKUP VALIDATE to tape is significantly less than the time for a real backup to tape, then writing to the output device is the likely bottleneck. See "Tuning the Copy and Write Phases".

https://docs.oracle.com/cd/E28271_01/backup.1111/e10642/rcmtunin.htm

C&D correct

upvoted 1 times

janw 4 years, 6 months ago

B & F can't be both right, f is incorrect, as the number of I/O processes is set to 4 only if asynch io is disabled, so only in the case of synch io, it can help ==> b correct, f incorrect

upvoted 1 times

CyborgQ 4 years, 6 months ago

Sorry: CEF is correct With synchronous I/O, it is difficult to identify specific bottlenecks because all synchronous I/O is a bottleneck to the process. The only way to tune synchronous I/O is to compare the rate (in bytes/second) with the device's maximum throughput rate. If the rate is lower than the rate that the device specifies, then consider tuning this aspect of the backup and restore process.

To determine the rate of synchronous I/O:

Start SQL*Plus and connect to the target database.

Query the DISCRETE_BYTES_PER_SECOND column in the V$BACKUP_SYNC_IO view to display the I/O rate.

If you see data in V$BACKUP_SYNC_IO, then the problem is that you have not enabled asynchronous I/O or you are not using disk I/O slaves.

upvoted 1 times

Question #34 *Topic 1*

You are managing this configuration:
1. CDB1 is a container database.
2. PDB1 and PDB2 are two pluggable databases in CDB1.
3. USER1.EMP is a table in PDB1 and USER2.DEPT is a table in PDB2.
CDB1 user SYS executes these commands after connecting successfully to PDB2:
Which two are true? (Choose two.)

```
SQL> ALTER SESSION SET CONTAINER=pdb1;
Session altered.

SQL> INSERT INTO user1.emp VALUES(100, 'Alan',1);
1 row created.

SQL> INSERT INTO user1.emp VALUES(101. 'Ben', 1);
1 row created.

SQL> ALTER SESSION SET CONTAINER=pdb2;
Session altered.

SQL> INSERT INTO user2.dept VALUES(1,'IT');
```

A. The inserts on USER1.EMP remain uncommitted when the session connected to PDB2.

B. The inserts on USER1.EMP were committed when the session inserted a row into USER2.DEPT.

C. The insert on USER2.DEPT fails because of the active transaction in the parent container.

D. The insert on USER2.DEPT is a recursive autonomous transaction by the child session and is committed.

E. The inserts on USER1.EMP were rolled back when the session connected to PDB2.

F. The insert on USER2.DEPT is uncommitted.

G. The inserts on USER1.EMP were committed when the session connected to PDB2.

Suggested Answer: *D, E*
Official Answer: DE

Community vote distribution
AC (100%)

you1234 Highly Voted 6 years, 1 month ago
other session to check any transaction

Version 19.5.0.0.0

SQL> SQL> alter session set container=ELC20MIG;

Session altered.

SQL> select * from user1.emp;

no rows selected

correct answer is A & C
upvoted 25 times

you1234 Highly Voted 6 years, 1 month ago
tested in 19c.

correct answer is A & C
upvoted 17 times

df404e2 Most Recent 1 year, 9 months ago

Selected Answer: AC

Instead of my previous comment, F is FALSE. This is because the insert failed, there is no uncommitted insert. There just is no insert at all.
upvoted 1 times

df404e2 1 year, 9 months ago

Answer: A, F F is TRUE, there is no commit at all. Reasoning for C being false might be the keyword "parent container".
upvoted 2 times

cujar2003 1 year, 9 months ago

Selected Answer: AC

Connected to: Oracle Database 19c Enterprise Edition Release 19.0.0.0.0 - Production Version 19.9.0.0.0

SQL> alter session set container=PDBDEV;

Session altered.

SQL> insert into hr.emp values(100,'acujar');

1 row created.

SQL> insert into hr.emp values(101,'arozo');

1 row created.

SQL> alter session set container=PDBDEV2;

Session altered.

SQL> insert into hr.dept values (100, 'afaa'); insert into hr.dept values (100, 'afaa') * ERROR at line 1: ORA-65023: active transaction exists in container PDBDEV

SQL>
upvoted 1 times

guretto 2 years, 8 months ago

A C From MyOracleSupport note "Oracle Multitenant: Frequently Asked Questions (Doc ID 1511619.1)"

Can a transaction span across PDBS ? No, though "alter session set container" is allowed after starting a transaction in a PDB, only select is allowed in the second PDB.Transaction is preserved and you can do commit or rollback after switch back to original PDB
upvoted 2 times

Shunpin 2 years, 8 months ago

Selected Answer: AC

SQL> alter session set container=ORCLPDB;

Session altered.

SQL> create table a (col1 varchar2(10));

Table created.

SQL> insert into a values ('1');

1 row created.

SQL> alter session set container=APPCON1;

Session altered.

SQL> create table a (col1 varchar2(10)); create table a (col1 varchar2(10)) * ERROR at line 1: ORA-65023: active transaction exists in container ORCLPDB

upvoted 2 times

Guhborges 3 years ago

Selected Answer: AC

A C 100%

Here my own test.

SQL> insert into in1 values(1);

1 row created.

SQL> alter session set container=pdb2;

Session altered.

SQL> insert into in2 values(2); insert into in2 values(2) * ERROR at line 1: ORA-65023: active transaction exists in container PDB1

upvoted 1 times

kaka321 3 years, 1 month ago

I will go for GF. all ddl command when issued on a session automatically commit any existing dml on that same session. that makes G correct.

upvoted 2 times

Guhborges 3 years ago

A and C 100%

upvoted 1 times

ScottL 3 years, 3 months ago

Selected Answer: AC

A and C

upvoted 1 times

gio 3 years, 3 months ago

Selected Answer: AC

i think ac

upvoted 1 times

vkra 3 years, 6 months ago

Selected Answer: AC

A+C is correct

upvoted 1 times

AliNawab 4 years, 11 months ago

Surely A & C sqlplus / as sysdba

SQL*Plus: Release 19.0.0.0.0 - Production on Tue Sep 14 16:58:07 2021 Version 19.11.0.0.0

Copyright (c) 1982, 2020, Oracle. All rights reserved.

Connected to: Oracle Database 19c Enterprise Edition Release 19.0.0.0.0 - Production Version 19.11.0.0.0

SQL> ALTER SESSION SET CONTAINER=PDB1;

Session altered.

SQL> INSERT INTO USER1.EMP VALUES(100);

1 row created.

SQL> INSERT INTO USER1.EMP VALUES(200);

1 row created.

SQL> ALTER SESSION SET CONTAINER=PDB2;

Session altered.

SQL> INSERT INTO USER2.DEPT VALUES(300); INSERT INTO USER2.DEPT VALUES(300) * ERROR at line 1: ORA-65023: active transaction exists in container PDB1

SQL>

upvoted 9 times

kameni 4 years, 11 months ago

A transaction cannot span multiple containers. If you start a transaction and use ALTER SESSION SET CONTAINER to switch to a different container, then you cannot issue DML, DDL, COMMIT, or ROLLBACK statements until you switch back to the container in which you started the transaction.

upvoted 2 times

Neil107 5 years ago

A and F. Refer to The 12c DBA Guide p. 399.

upvoted 1 times

ObserverPL 5 years, 3 months ago

SQL> alter session set container=PDB1;

Session altered.

SQL> insert into emp values(1);

1 row created.

SQL> alter session set container=PDB2;

Session altered.

SQL> insert into emp values(2); insert into emp values(2) * ERROR at line 1: ORA-65023: active transaction exists in container PDB1

upvoted 4 times

you1234 6 years, 1 month ago

SQL> alter session set container=ELC20MIG;

Session altered.

SQL> select * from user1.emp;

no rows selected

SQL> insert into user1.emp values(100);

1 row created.

SQL> insert into user1.emp values(200);

1 row created.

SQL> alter session set container=PDB2;

Session altered.

SQL> insert into user2.emp values(100); insert into user2.emp values(100) * ERROR at line 1: ORA-65023: active transaction exists in container ELC20MIG

upvoted 6 times

Question #35 *Topic 1*

Examine this configuration:
1. CDB1 is an Oracle Database 12c Release 2 database containing pluggable databases PDB$SEED, PDB1, and PDB2.
2. PDB$SEED is open READ ONLY
3. PDB1 is open READ WRITE
4. PDB2 is MOUNTED.
5. ORACLE_HOME is /u01/app/oracle/product/18.1.0/dbhome_1.
You execute these commands before upgrading the database to the current release:
For which databases will fixup scripts be created?

```
$ . oraenv
ORACLE_SID = [cdb1] ? cdb1
The Oracle base remains unchanged with value /u01/app/oracle

$ $ORACLE_HOME/jdk/bin/java -jar preupgrade.jar TERMINAL TEXT
```

A. CDB1, PDB$SEED, PDB1, and PDB2

B. PDB$SEED, PDB1, and PDB2 only

C. CDB1 and PDB$SEED only

D. CDB1, PDB1, and PDB2 only

E. CDB1, PDB$SEED, and PDB1 only

Suggested Answer: *E*
Official Answer: E

Community vote distribution
E (100%)

janw Highly Voted 5 years ago
e correct!

For multitenant architecture (CDB and PDB) upgrades, open up all the PDBs that you want the tool to analyze before you run the tool.
upvoted 12 times

Alejandrrro 4 years, 6 months ago
Agree with you.

Also tested in my database:

[oracle@apex tmp]$ $ORA-CLE_HOME/jdk/bin/java -jar preupgrade.jar TERMINAL TEXT INFORMATION - Container: XEPDB2 is in: MOUNTED status and will not be processed unless opened.
upvoted 1 times

ObserverPL Highly Voted 4 years, 3 months ago
I agree, E
upvoted 5 times

Alvinzzz 2 years, 1 month ago
I think all of them are wrong as PDB$SEED is not open for generate access or modification, preupgrade.jar cannot directly generate fixup script for it.

But from the answers above, E is the most correct answer although it includes PDB$SEED. A, B & D must be wrong as PDB2 is wrong.
upvoted 1 times

Guhborges Most Recent 2 years ago
E correct.

The script will not perform a change, but will fetch information.

upvoted 1 times

ScottL 2 years, 3 months ago

Selected Answer: E

E correct

upvoted 1 times

gio 2 years, 3 months ago

Selected Answer: E

i think e

upvoted 1 times

vkra 2 years, 6 months ago

Selected Answer: E

E is correct

upvoted 1 times

marcinb32 5 years, 1 month ago

E is correct https://docs.oracle.com/en/database/oracle/oracle-database/18/spmsu/running-pre-upgrade-information-tool-for-non-cdb-checks.html#GUID-B04CDE73-45F1-4717-98BB-9FCC209DEC01

upvoted 5 times

veginha 5 years, 1 month ago

i think E is correct, right?

upvoted 4 times

Question #36 *Topic 1*

Which two are true about Oracle Flashback features? (Choose two.)

A. FLASHBACK QUERY can retrieve REDO records from ONLINE and ARCHIVED REDO LOG files.

B. FLASHBACK VERSION QUERY can retrieve REDO records from ONLINE and ARCHIVED REDO LOG files.

C. FLASHBACK TABLE can undrop a column.

D. FLASHBACK DROP can undrop an index when undropping a table.

E. After a database is restored from flashback logs using the FLASHBACK DATABASE command, it is sometimes rolled forward using redo logs.

Suggested Answer: *A, B*
Official Answer: AB

monad2006 Highly Voted 4 years, 8 months ago
not a-b because flashback use undo not redo/archive. not c because flashback table cannot rever ddl d-e -> as of before drop retrieve from the recycle bin also the related index, and after using flashback database command, sometimes you could go forward with redo.
upvoted 21 times

veginha Highly Voted 4 years, 9 months ago
This response is wrong, flashback query use UNDO, not archivelogs, i think right answer must be D and E
upvoted 16 times

ScottL Most Recent 1 year, 9 months ago
Selected Answer: DE
D,E correct
upvoted 2 times

gio 1 year, 9 months ago
Selected Answer: DE
i think de
upvoted 1 times

vkra 2 years ago
Selected Answer: DE
D + E is correct
upvoted 1 times

hilaire 2 years, 1 month ago
D E correct
upvoted 2 times

Abdou1968 4 years, 4 months ago
D & E are Correct
upvoted 10 times

Abdou1968 4 years, 4 months ago
A, B are Wrong. Flashback Query & Version Query need Undo not Redo C is Wrong. D is Correct. Assuming that an index is not dropped before dropping a table. E is Correct
upvoted 8 times

pgtips 4 years, 6 months ago
Agree with D & E for the same reason. all other uses UNDO.
upvoted 8 times

janw 4 years, 6 months ago
d,e correct!
upvoted 8 times

Question #37 *Topic 1*

Which three are true about an application seed pluggable database (PDB)? (Choose three.)

A. It is automatically synchronized with its application root PDB when an application is upgraded.

B. It cannot be added to an application container after the application container has already been created.

C. A new application PDB created by cloning an application seed PDB can have an old version of the application installed after cloning completes.

D. It is automatically synchronized with its application root PDB when an application is installed.

E. It cannot be dropped from its application container.

F. A new application PDB created by cloning an application seed PDB can have an up-to-date version of the application installed after cloning completes.

G. It is not required in an application container.

Suggested Answer: *D, F, G*
Official Answer: DFG

Community vote distribution
CFG (100%)

janw Highly Voted 6 years ago
cfg: agree
upvoted 15 times

marcinb32 Highly Voted 6 years, 1 month ago
CFG ??
upvoted 6 times

taotsumiau 6 years, 1 month ago
Agree CFG https://docs.oracle.com/en/database/oracle/oracle-database/19/multi/creating-removing-application-containers-seeds-with-sql-plus.html#GUID-065649AC-18ED-4858-ACC5-36F011362A83
upvoted 5 times

cujar2003 Most Recent 1 year, 9 months ago
Selected Answer: CFG
Agree CFG
upvoted 1 times

gio 3 years ago
Selected Answer: CFG
CFG 100%
upvoted 1 times

Blob44 3 years, 1 month ago
Selected Answer: CFG
- An application seed is an optional, user-created PDB within an application container. You can use the DROP PLUGGABLE DATABASE statement to drop an application seed. You can drop an application seed when you no longer need it.

- After the application seed is created, it is synchronized with the application root so that the applications are installed in the application seed.

- When an application in the application root is upgraded or patched, the application seed must be synchronized with the application root to apply these changes.
upvoted 1 times

Blob44 3 years, 1 month ago
https://docs.oracle.com/en/database/oracle/oracle-database/21/multi/creating-application-seeds.html#GUID-C1D1CCE5-5259-4527-A3F9-4E9DDCD61B06
upvoted 1 times

ScottL 3 years, 3 months ago

Selected Answer: CFG

CFG correct

upvoted 1 times

vkra 3 years, 6 months ago

Selected Answer: CFG

CFG is correct

upvoted 2 times

Dhanushbh 5 years, 1 month ago

To create an application container, the current container must be the CDB root and you must specify the AS APPLICATION CONTAINER clause in the CREATE PLUGGABLE DATABASE statement. You must create the application container using Oracle Managed Files. so F is wrong

Note: An application container cannot be unplugged or dropped if any application PDBs belong to it. so E is correct

upvoted 1 times

xRodge 5 years ago

@Dhanushbh : I believe that you may have read the question incorrectly. The question is asking about what is true about an "application SEED pluggable database". Answer E is not saying that you drop the application container, merely that you drop the SEED pluggable database. This is perfectly allowable, oracle-base displays this exact thing happening here: https://oracle-base.com/articles/12c/multitenant-application-containers-12cr2

We can drop an existing application seed like any other PDB. This has no impact on application container PDBs created using it.

CONN / AS SYSDBA ALTER SESSION SET container = appcon1;

ALTER PLUGGABLE DATABASE appcon1$SEED CLOSE; DROP PLUGGABLE DATABASE appcon1$SEED INCLUDING DATAFILES;

E is, definitely, incorrect.

upvoted 2 times

DBA4All 4 years, 2 months ago

Dropping an Application Seed: https://docs.oracle.com/en/database/oracle/oracle-database/19/multi/creating-removing-application-containers-seeds-with-sql-plus.html#GUID-93AAE6F0-7969-4D05-B10B-F2514CB39417

upvoted 1 times

Question #38 *Topic 1*

Examine this configuration:
1. CDB1 is a container database.
2. PDB1 and PDB2 are pluggable databases in CDB1.
3. PDB1 and PDB2 are OPEN in READ WRITE mode.
You execute these commands successfully:
Which two are true? (Choose two.)

```
$ export ORACLE_SID=CDB1
$ sqlplus / as sysdba

SQL> ALTER SESSION SET CONTAINER = PDB1;
Session altered.

SQL> SHUTDOWN IMMEDIATE
```

A. Uncommitted transactions in PDB1 have been rolled back.

B. PDB1 is closed.

C. Uncommitted transactions in CDB1 and PDB1 have been rolled back.

D. CDB1 is shut down.

E. CDB1 is in MOUNT state

Suggested Answer: *B, E*
Official Answer: BE

taotsumiau Highly Voted 6 years, 2 months ago
I had test this on my lab. the Answer should be AB
upvoted 25 times

you1234 Highly Voted 6 years, 1 month ago
A & B is correct
upvoted 12 times

cujar2003 Most Recent 1 year, 9 months ago
Selected Answer: AB
SQL> alter session set container=PDBDEV;

Session altered.

SQL> shutdown immediate; ORA-01097: cannot shutdown while in a transaction - commit or rollback first

SQL> rollback;

Rollback complete.

SQL> shutdown immediate; Pluggable Database closed. SQL> exit
upvoted 1 times

pitt84 1 year, 10 months ago
AE PDB in mount state
upvoted 1 times

gedde 2 years, 5 months ago
With 19c it does not let you shutdown immediate and says : SQL> shutdown immediate ORA-01097: cannot shutdown while in a transaction - commit or rollback first
upvoted 1 times

kaka321 3 years, 1 month ago

AB is absolutely correct

upvoted 1 times

Guhborges 3 years, 1 month ago

Selected Answer: AB

AB 110%

upvoted 1 times

Blob44 3 years, 1 month ago

Selected Answer: AB

https://oracle-base.com/articles/12c/multitenant-connecting-to-cdb-and-pdb-12cr1

upvoted 1 times

ScottL 3 years, 3 months ago

Selected Answer: AB

AB correct

upvoted 1 times

gio 3 years, 3 months ago

Selected Answer: AB

i think ab

upvoted 1 times

jackymak 3 years, 9 months ago

Selected Answer: AB

https://www.oracletutorial.com/oracle-administration/oracle-shutdown/ During the SHUTDOWN IMMEDIATE, all the connected sessions are disconnected immediately, all uncommitted transactions are rolled back, and the database completely shuts down.

upvoted 2 times

Trololoshko 3 years, 10 months ago

A & B 100500 %

upvoted 3 times

nobody347 5 years ago

A & B 100%

upvoted 5 times

pgtips 6 years ago

Should be A & B

upvoted 8 times

janw 6 years ago

a,b: ok

upvoted 7 times

CyborgQ 6 years ago

A,B is only correct answers

upvoted 8 times

Question #39 *Topic 1*

Which three are true about Automatic Workload Repository (AWR), Automatic Database Diagnostic Monitor (ADDM), and the Manageability Monitor (MMON) background process? (Choose three.)

A. ADDM can recommend shrinking the buffer cache.

B. ADDM can recommend extending the buffer cache.

C. By default, MMON creates an AWR snapshot every 30 minutes.

D. ADDM performs its analysis only when a DBA requests it.

E. By default, AWR snapshots are automatically purged after eight days.

F. AWR snapshots must be deleted when no longer required by ADDM.

Suggested Answer: *A, E, F*
Reference: https://books.google.com.pk/books?id=DeQkU5E75DUC&pg=PA869&lpg=PA869&dq=By+default,+AWR+snapshots+are+automatically+purged+after +eight +days&source=bl&ots=DfoGmJfCMy&sig=ACfU3U0Pf2RgIs72LUV4HUzs5UIXI8sXEQ&hl=en&sa=X&ved=2ahUKEwjb8bKhu9noAhUwTxUIHRpYAuMQ6AEwA3o
ECAsQLA#v=onepage&q=By%20default%2C%20AWR%20snapshots%20are%20automatically%20purged%20after%20eight%20days&f=false

Community vote distribution



ald85 Highly Voted 5 years, 2 months ago
ABE F incorrect because AWR snapshots are deleted automatically.
upvoted 23 times

ama 4 years, 6 months ago
AWR Snapshot's default interval is 1 hour and retention is 8 days. It means All AWR statistics are collected every 1 hour and they are kept 8 days by default.

That means E is correct and F is wrong… so final correct answers are A; B; E
upvoted 5 times

janw Highly Voted 5 years ago
agree with abe
upvoted 11 times

Guhborges Most Recent 2 years, 1 month ago
Selected Answer: ABE
A,B,E i think
upvoted 1 times

Dhanushbh 2 years, 2 months ago
AWR (Automatic Workload Repository) snapshots cannot be created directly in the CDB$ROOT container in Oracle Database. AWR is a feature of Oracle Database that collects and stores performance statistics for diagnosing performance issues.

B is a wrong answer
upvoted 1 times

ScottL 2 years, 3 months ago
ABE correct
upvoted 1 times

gio 2 years, 3 months ago
Selected Answer: ABE
i think abe
upvoted 1 times

Yennism 2 years, 8 months ago

ABE. F =false. AWR snapshots must be deleted when no longer required by ADDM. => This is automatically

upvoted 2 times

flaviogcmelo 2 years, 11 months ago

Selected Answer: ABE

ABE Ref.: https://docs.oracle.com/en/database/oracle/oracle-database/19/tgdba/automatic-performance-diagnostics.html#GUID-0C89C9A5-A9EE-4E65-95F8-208C9EA36743 Ref.: https://docs.oracle.com/en/database/oracle/oracle-database/19/tgdba/gathering-database-statistics.html#GUID-144711F9-85AE-4281-B548-3E01280F9A56

upvoted 2 times

rashedptdba 3 years, 10 months ago

Corrent:A,B,E

upvoted 5 times

Neil107 4 years ago

It's ABE.

upvoted 5 times

undoundo 4 years, 9 months ago

Is F the correct answer? Or is X correct? B...□□C...×□D...×□E...□

upvoted 1 times

undoundo 4 years, 9 months ago

↑ I made a mistake. Is F the correct answer? Or is A correct? B...□□C...×□D...×□E...□

upvoted 1 times

cerebro2000x 5 years ago

abe snapshots that are 8 days old are auto purge

upvoted 7 times

taotsumiau 5 years, 2 months ago

BEF ??

upvoted 2 times

Question #40 *Topic 1*

Examine the command for creating pluggable database PDB2 in container database CDB2.
Select three options, any one of which is required for it to execute successfully. (Choose three.)

```
CREATE PLUGGABLE DATABASE pdb2
       ADMIN USER pdb2_adm
       IDENTIFIED BY 123pdb
       ROLES=(CONNECT);
```

A. Add the FILE_NAME_CONVERT clause to the statement and set the PDB_FILE_NAME_CONVERT parameter.

B. Add only the CREATE_FILE_DEST clause to the statement.

C. Set only the PDB_FILE_NAME_CONVERT parameter.

D. Set the PDB_FILE_NAME_CONVERT parameter and enable OMF.

E. Enable only OMF.

F. Add the FILE_NAME_CONVERT clause to the statement and enable Oracle Managed Files (OMF)

Suggested Answer: *B, D, E*
Official Answer: BDE

Community vote distribution
BCE (100%)

veginha Highly Voted 5 years, 8 months ago
All are valid, but i think BCE is correct because you don't require the file_name_convert in option "A" just pdb_file_name_convert, don't need pdb_file_name_convert in "D" just omf, also don't require file_name_convert in option "F" just omf. Any comment?
upvoted 15 times

mihanik Highly Voted 5 years, 6 months ago
BCE https://docs.oracle.com/en/database/oracle/oracle-database/18/multi/overview-of-pdb-creation.html#GUID-39D2F826-7460-4439-B410-29E16C0CA0B1 not A PDB_FILE_NAME_CONVERT maps names of existing files to new file names when processing a CREATE PLUGGABLE DATABASE statement, as well as when processing the ENABLE PLUGGABLE DATABASE clause of the CREATE DATABASE statement, if the file_name_convert_clause is not specified and Oracle Managed Files is not enabled. B Enough to create OMF files C Looks good too not D Coz of using OMF not requires PDB_FILE_NAME_CONVERT E DB_CREATE_FILE_DEST? Thats enough to create PDB not F Coz of using OMF not requires FILE_NAME_CONVERT
upvoted 10 times

guretto Most Recent 2 years, 2 months ago
BCE For https://smarttechways.com/2019/03/05/ora-65016-file_name_convert-must-be-specified/

B C E are true F is false

For https://docs.oracle.com/en/database/oracle/oracle-database/19/refrn/PDB_FILE_NAME_CONVERT.html#GUID-074F8896-D565-4139-BCDB-C81C9D741941): A D are false
upvoted 1 times

Blob44 2 years, 7 months ago
Selected Answer: BCE
see mihanik explanation
upvoted 1 times

gio 2 years, 8 months ago
Selected Answer: BCE
source: http://dbaparadise.com/2021/03/file_name_convert-vs-pdb_file_name_convert/
upvoted 2 times

ScottL 2 years, 9 months ago
Selected Answer: BCE

BCE Correct https://oracle-base.com/articles/12c/multitenant-create-and-configure-pluggable-database-12cr1

upvoted 2 times

vkra 3 years ago

Selected Answer: BCE

BCE is correct

upvoted 1 times

Neil107 4 years, 6 months ago

BCE sounds right.

upvoted 5 times

RinD 4 years, 11 months ago

Agree with BCE

Reference: https://docs.oracle.com/database/121/SQLRF/statements_6010.htm#SQLRF55686

You can specify FILE_NAME_CONVERT = NONE, which is the same as omitting this clause. If you omit this clause, then the database first attempts to use Oracle Managed Files to generate file names. If you are not using Oracle Managed Files, then the database uses the PDB_FILE_NAME_CONVERT initialization parameter to generate file names. If this parameter is not set, then an error occurs.

upvoted 3 times

janw 5 years, 6 months ago

BCE is correct

upvoted 6 times

nyanyanyao0826 5 years, 7 months ago

Correct = ACE B = CREATE~ must enable OMF D,F = OMFdon't enable

upvoted 3 times

Das97 3 years, 1 month ago

B is correct: This CREATE_FILE_DEST lets you override the default behaviour. You can enable or disable Oracle Managed Files for the PDB and you specify a different base file system directory or Oracle ASM disk group for the PDB's files. If you specify a value other than NONE, then the database implicitly sets the DB_CREATE_FILE_DEST initialization parameter with SCOPE=SPFILE in the PDB.
(https://docs.oracle.com/database/121/SQLRF/statements_6010.htm#CACFDEHC)

upvoted 1 times

julica 5 years, 7 months ago

I think only E is right:

Version 19.6.0.0.0

SQL> show parameter CREATE_FILE

NAME TYPE VALUE ---------- ---------- ---------- db_create_file_dest string +DATA

SQL> show parameter file_name_convert

NAME TYPE VALUE ---------- ---------- ---------- db_file_name_convert string log_file_name_convert string pdb_file_name_convert string SQL>

SQL> create pluggable database pdb3 admin user pdb3_adm identified by welcome1 roles=(connect);

Pluggable database created.

SQL> SQL> show pdbs

CON_ID CON_NAME OPEN MODE RESTRICTED --------- ---------- --------- --------- 2 PDB$SEED READ ONLY YES 3 PDB1 READ WRITE YES 4 PDB2 READ WRITE YES 5 PDB3 MOUNTED SQL>

upvoted 1 times

Question #41 *Topic 1*

Which two are true about Recovery Manager (RMAN) diagnostic message output? (Choose two.)

A. Media Management messages for SBT devices are always written to sbtio.log.

B. RMAN error stacks should be read from the bottom up as that is the order in which errors are generated.

C. RMAN error stacks should be read from the top down as that is the order in which errors are generated.

D. The RMAN LOG command line clause causes output issued during RMAN command compilation to be written to a log file and to standard output.

E. The RMAN LOG command line clause causes output issued during RMAN command compilation to be written to a log file only.

F. Media Management messages for SBT devices are written to an Oracle trace file.

Suggested Answer: *E, F*
Official Answer: EF

Community vote distribution
BE (100%)

julica Highly Voted 6 years, 1 month ago
Correct are B and E:

[oracle@rac1 admin]$ rman target / log '/tmp/log.txt' RMAN> list backup summary; RMAN> exit [oracle@rac1 admin]$ cat /tmp/log.txt

Recovery Manager: Release 19.0.0.0.0 - Production on Mon Jun 22 15:48:39 2020 Version 19.6.0.0.0

Copyright (c) 1982, 2019, Oracle and/or its affiliates. All rights reserved.

connected to target database: TST193 (DBID=1785824688)

RMAN> using target database control file instead of recovery catalog

List of Backups ========== Key TY LV S Device Type Completion Time #Pieces #Copies Compressed Tag ------- -- -- - ---------- --------- ------- ------- ---------- --- 1 B F A DISK 15-JUN-20 1 1 NO TAG20200615T165352 2 B F A DISK 15-JUN-20 1 1 NO TAG20200615T174406 3 B F A DISK 22-JUN-20 1 1 NO TAG20200622T095655 4 B F A DISK 22-JUN-20 1 1 NO TAG20200622T095942 5 B F A DISK 22-JUN-20 1 1 NO TAG20200622T154703

RMAN>

Recovery Manager complete.
upvoted 21 times

Alejandrrro 5 years, 6 months ago
Agree with B,E

Here's the an additional reference

https://docs.oracle.com/database/121/BRADV/rcmtroub.htm#BRADV89875

The sbtio.log contains information written by the media management software, not Oracle Database. Thus, you must consult your media vendor documentation to interpret the error codes and messages. If no information is written to the sbtio.log, then contact your media manager support to ask whether they are writing error messages in some other location, or whether there are steps you must take to have the media manager errors appear in sbtio.log.

Read the messages from the bottom up, because this is the order in which RMAN issues the messages. The last one or two errors displayed in the stack are often the most informative.

upvoted 5 times

janw Highly Voted 6 years ago

agree with B,E

upvoted 8 times

cujar2003 Most Recent 1 year, 9 months ago

Selected Answer: BE

https://docs.oracle.com/database/121/BRADV/rcmtroub.htm#BRADV89875

upvoted 1 times

ScottL 3 years, 3 months ago

Selected Answer: BE

BE correct

upvoted 1 times

gio 3 years, 3 months ago

Selected Answer: BE

i think be: https://docs.oracle.com/database/121/BRADV/rcmtroub.htm#GUID-3F0A6F9D-1FD0-4507-8E2B-B1A205690BCC

upvoted 1 times

jareach 3 years, 6 months ago

Selected Answer: BE

RMAN error stacks goes from general to specific, or from consequence to a cause when read top-down. The root cause is always on the bottom.

upvoted 2 times

Yennism 3 years, 8 months ago

BE but sound that A is true too.

upvoted 1 times

guretto 2 years, 8 months ago

It's software vendor dependent: "If no information is written to the sbtio.log, then contact your media manager support to ask whether they are writing error messages in some other location" from https://docs.oracle.com/database/121/BRADV/rcmtroub.htm#BRADV89875

upvoted 1 times

ErikJanssen 3 years, 9 months ago

Option A looks valid too. https://docs.oracle.com/database/121/BRADV/rcmtroub.htm#BRADV012

upvoted 1 times

Neil107 5 years ago

I think it's BE.

upvoted 3 times

asefa 5 years, 6 months ago

the correct answer is B and E because RMAN error messages read from bottom up and log option writes output to log files only.

upvoted 4 times

you1234 6 years, 1 month ago

B & D is correct

upvoted 2 times

logica 6 years, 1 month ago

I think B, D and F are correct

upvoted 1 times

ald85 6 years, 2 months ago

I think BD. B sure that is correct, the other i doubt between D or F, any suggestion?

upvoted 1 times

logica 6 years, 1 month ago

D is incorrect, Rman's LOG parameter cause the output to be written to a log only

upvoted 3 times

ald85 6 years, 1 month ago

Finally BE. https://logic.edchen.org/3-ways-to-set-rman-log-location/

upvoted 8 times

Question #42 *Topic 1*

Which three are true about managing memory components in an Oracle database instance? (Choose three.)

A. With Automatic Shared Memory Management, the database instance can increase the Large Pool size by reducing the Shared Pool size.

B. With Automatic Memory Management, the database instance can increase the System Global Area size by reducing the Program Global Area size.

C. Automatically tuned and resized System Global Area components will always revert to their initial sizes after an instance restart.

D. Automatic Memory Management must be used together with locking the System Global Area into physical memory.

E. With Automatic Shared Memory Management, the database instance can increase the Program Global Area size by reducing the System Global Area size.

F. On Line Transaction Processing systems often use less Program Global Area than Decision Support Systems.

Suggested Answer: *A, E, F*

Reference: https://docs.oracle.com/database/121/TGDBA/tune_shared_pool.htm

Community vote distribution

ABF (100%)

veginha Highly Voted 4 years, 8 months ago

ABF in AMM the instance exchanges memory between the SGA and the instance PGA as needed (B is correct) In ASMM the instance automatically distributes this memory among the various SGA components (JUST SGA) (A is correct, E incorrect) After a restart the instance return to the SGA values before the shutdown (not the initial values) (C incorrect) https://docs.oracle.com/en/database/oracle/oracle-database/19/admin/managing-memory.html#GUID-8F54391B-D42A-4FDA-9D12-E1F81FD113EA

F is also correct: For OLTP systems, the PGA memory typically makes up a small fraction of the available memory, leaving most of the remaining memory for the SGA. For DSS systems running large, memory-intensive queries, PGA memory can typically use up to 70% of the available memory. https://docs.oracle.com/database/121/TGDBA/tune_pga.htm#TGDBA472

upvoted 22 times

zubi_7 Highly Voted 4 years, 5 months ago

ABF is the correct answer

upvoted 9 times

ScottL Most Recent 1 year, 9 months ago

Selected Answer: ABF

ABF correct

upvoted 1 times

gio 1 year, 9 months ago

Selected Answer: ABF

i think abf

upvoted 1 times

jareach 2 years ago

Selected Answer: ABF

E is incorrect as the PGA is not a part of SGA. So the Shared Memory Management cannot mess up with PGA.

upvoted 1 times

Jibss 2 years, 5 months ago

The correct answer is ACF. The SGA components memory are allocated at instance startup you can set the sga_target and sga_max_size to determine range allocated but distribution within the components is done at instance startup. B is not correct because SGA and PGA are set differently using SGA_TARGET & SGA_MAX_SIZE and PGA is set with PGA_AGGREGATE_TARGET & PGA_AGGREGATE_LIMIT. SGA is a server process memory area while PGA is a user process memory area.

upvoted 1 times

apagnini 4 years, 5 months ago

Automatic Memory Management (AMM) is enabled with "MEMORY_TARGET" or/and "MEMORY_MAX_TARGET" init parameters. Dynamic components of Automatic Memory Management (AMM) are:

SELECT component, current_size, min_size, max_size FROM v$memory_dynamic_components WHERE current_size != 0;

COMPONENT CURRENT_SIZE MIN_SIZE MAX_SIZE ---------- ---------- ---------- ---------- shared pool 197132288 192937984 197132288 large pool 4194304 4194304 4194304 java pool 41943040 41943040 41943040 SGA Target 318767104 285212672 318767104 DEFAULT buffer cache 71303168 41943040 75497472 PGA Target 104857600 104857600 138412032

6 rows selected.

SQL> then I think the correct answer is ABE

upvoted 1 times

SimoneF 3 years, 9 months ago

Careful, E states "Automatic SHARED Memory Management", so it's ASMM, not AMM. ASMM can only manage the distribution of memory inside the SGA, not between it and PGA.

upvoted 3 times

cerebro2000x 4 years, 6 months ago

abf for me

upvoted 2 times

cerebro2000x 4 years, 6 months ago

c is a wrong answer bec values are recorded to spfile so after a restart no need to relearn their values

upvoted 2 times

cerebro2000x 4 years, 6 months ago

d is a wrong answer bec of the word MUST, servers using huge pages or pagefiles cant use amm so these servers are only into asmm

upvoted 1 times

julica 4 years, 7 months ago

ABF is correct

upvoted 3 times

Question #43 *Topic 1*

A database is configured in ARCHIVELOG mode.
Full RMAN backups are taken daily and no backup to trace has been taken of the control file.
A media failure has occurred.
In which two scenarios is complete recovery possible? (Choose two.)

A. when any archived log from, before, or after the most recent backup is corrupt.

B. after losing all copies of the control file

C. after losing an archived log from after the most recent backup

D. after losing an archived log from before the most recent backup

E. after losing the SYSTEM tablespace

Suggested Answer: *A, D*
Official Answer: AD

Community vote distribution
DE (100%)

cerebro2000x Highly Voted 5 years, 6 months ago
de b will do a resetlogs
upvoted 22 times

1451755 Most Recent 1 year, 4 months ago
I vote for DE, but disagree that b will do a resetlogs.

If there is no backup of controlfile in trace and backup of controlfile at all, no copy of controlfile, DBA may still know all information to re-create controlfile !!! (values like MAXLOGFILES 16, MAXLOGMEMBERS 3, MAXDATAFILES 100, ...) since lot of databases uses dbca defaults or have other documentation. Then refer section "Set #1. NORESETLOGS" in https://prakashoracledba.wordpress.com/2015/02/16/step-by-step-alter-database-backup-controlfile-to-trace/ ... I did sucessfully fixed this way 19c db. with lost of ontrol file (disk error) and no backup of contfolfile at all, no copy (db. was in archived mode). I'm living proof that B is also right, but it's a bit hacking.
upvoted 1 times

veginha Highly Voted 5 years, 8 months ago
i think D and E are correct. If you lose any archive from after the backup, you can't do a complete recovery, so A and C are discarted, also you can't mount the database if you don't have a valid control file. (not B)
upvoted 8 times

taotsumiau 5 years, 8 months ago
what if the questions is incomplete recovery, what do you think the answer? I think C,D are correct, what's your idea?
upvoted 1 times

veginha 5 years, 8 months ago
well if the question was "incomplete recovery" then i think that you only need the control file and the backup, so it could be possible to do the incomplete recovery in A,C,D, and also E, fortunately they ask for the complete recovery scenario, so i'm going with D, E :)
upvoted 3 times

dancymonkey 2 years ago
DE Complete Recovery Apply online or archived redo logs, or a combination of the two In this case, we need choice C to do complete recovery

Incomplete Recovery you do not apply all of the redo records generated after the most recent backup.

https://docs.oracle.com/cd/B13789_01/server.101/b10743/backrec.htm
upvoted 1 times

eleeitor 2 years, 7 months ago
B is a weird one. If autobackup is ON, we would have an extra backup piece for the control file in the backupset. (for CDBs and non-CDBs, starting with 12.0.0 or higher, autobackup is turned on) Even if autoback is OFF, then, when we issue a BACKUP DATABASE or BACKUP TABLESPACE that includes data file 1 (system), then by default oracle includes the control file as part of the backup, but not as a separate backup piece.

https://docs.oracle.com/en/database/oracle/oracle-database/19/rcmrf/BACKUP.html#GUID-73642FF2-43C5-48B2-9969-99001C52EB50

Why i would say B is wrong is because we somehow lost "ALL copies of the control files"

upvoted 1 times

Blob44 2 years, 7 months ago

Selected Answer: DE

Control File is not involved in media recovery Complete recovery needs backup + archivelogs to recover DB/TBS

upvoted 1 times

ScottL 2 years, 9 months ago

Selected Answer: DE

DE correct

upvoted 1 times

gio 2 years, 9 months ago

Selected Answer: DE

i think de

upvoted 1 times

jareach 3 years ago

Selected Answer: DE

All other situations will either deprive you of some transactional information }current or archived redo logs to redo the transactions up to date) or you do not have current LSN to complete the recovery (as the LSN is written in the controlfile header).

upvoted 1 times

hilaire 3 years, 1 month ago

DE correct sure

upvoted 1 times

G_C 3 years, 4 months ago

Selected Answer: DE

A - FALSE of course - database recovery is possible ony to SCN with, the last transaction in previous archivied redo log, B - TRUE/FALSE because we do not know how the backup was taken, C - FALSE (no any doubts), D - TRUE (no any doubts), E - TRUE (because we can restore full database),

upvoted 1 times

Alvinzzz 3 years, 8 months ago

You cannot open a DB when recoverd using backup controlfiles. RESETLOGS is required (in-complete recovery).

upvoted 2 times

abaisyet 4 years, 6 months ago

B and D

thegeekstuff.com/2014/11/oracle-rman-restore/ This is a complete recovery even though the database is opened with resetlogs option:

RMAN> STARTUP NOMOUNT; RMAN> RESTORE CONTROLFILE FROM "/backup/rman/ctl_c-12345-20141003-03"; RMAN> RESTORE DATABASE; RMAN> RECOVER DATABASE; RMAN> ALTER DATABASE OPEN RESETLOGS;

upvoted 2 times

G_C 3 years, 4 months ago

if only non default parameter CONFIGURE CONTROLFILE AUTOBACKUP ON or you have taken your backup using following command RUN { BACKUP DATABASE; BACKUP SPFILE; } # or if your database does not use a server parameter file: RUN { BACKUP DATABASE; BACKUP CURRENT CONTROLFILE; } but we do not know that ...

upvoted 1 times

dacoben415lywenw 5 years, 7 months ago

DE, also B is correct. Full backup includes controlfile even without CF copy so it can be restored from there...

upvoted 1 times

nww 5 years, 3 months ago

B is correct if CONFIGURE CONTROLFILE AUTOBACKUP is set to ON.

upvoted 5 times

Alvinzzz 2 years, 7 months ago

If controlfile need to restore from autobackup, it may not be the latest version. In such case, should it be in-complete recovery instead of complete recovery?

upvoted 1 times

Question #44 *Topic 1*

Which three are true about Database Point-in-Time Recovery? (Choose three.)

A. The database must have FLASHBACK DATABASE ON to perform Database Point-in-Time Recovery.

B. The database must be in MOUNT state when performing Database Point-in-Time Recovery.

C. Database Point-in-Time Recovery is performed by the Managed Recovery Process (MRP)

D. The Database must be in ARCHIVELOG mode.

E. The target point for the recovery must be specified as a stime or System Change Number (SCN).

F. The database must be open RESETLOGS after Database Point-in-Time Recovery.

Suggested Answer: *A, C, F*
Official Answer: ACF

Community vote distribution
BDF (100%)

veginha Highly Voted 5 years, 8 months ago
Response: B,D,F Archivelog mode and mount state are needed (B, and D) After the restore you need to open with reset logs (F) E is discarted because sequence number also is valid
https://docs.oracle.com/cd/B19306_01/backup.102/b14192/flashptr006.htm#:~:text=Database%20point%2Din%2Dtime%20recovery%20(DBPITR)%20restores%20the,forward%20to%20the%20target%20time.
upvoted 28 times

taotsumiau 5 years, 8 months ago
Agree, BDF
upvoted 6 times

dancymonkey Most Recent 2 years ago
Agreed, BEF A is wrong To perform Database Point-in-Time Recovery not only must have FLASHBACK DATABASE ON
upvoted 1 times

dancymonkey 2 years ago
Sorry I meant BDF
upvoted 1 times

janw Highly Voted 5 years, 6 months ago
agree BDF
upvoted 7 times

3c4de23 2 years, 1 month ago
MRP process is for a Dataguard recovery process. The FLASHBACK and PITR are separate features and independent of each other and hence A is out and D (You can recover a DB to a sequence number which is PITR) and E are incorrect as well Correct answer is BDF!
upvoted 1 times

Blob44 2 years, 7 months ago
Selected Answer: BDF
Certain prerequisites must be met to perform database point-in-time recovery (DBPITR). This includes the following:

Your database must be running in ARCHIVELOG mode.

You must have backups of all data files from before the target SCN for DBPITR and archived logs for the period between the SCN of the backups and the target SCN.

https://docs.oracle.com/en/database/oracle/oracle-database/19/bradv/rman-performing-flashback-dbpitr.html#GUID-2CF6EEFB-0344-4FD6-9BBB-F2D241172BC6
upvoted 1 times

ScottL 2 years, 9 months ago

Selected Answer: BDF

BDF correct

upvoted 1 times

gio 2 years, 9 months ago

Selected Answer: BDF

i think bdf

upvoted 1 times

jareach 3 years ago

Selected Answer: BDF

Must be in the MOUNT state, archivelog is needed to roll forward the transactions since the data files backup, and open resetlogs is needed because (not all) information in the online redo log is longer relevant.

upvoted 1 times

tlcecertified 5 years, 1 month ago

C is wrong, it's not in a data guard standby database environment

upvoted 1 times

cerebro2000x 5 years, 6 months ago

BDF

C can also use RESTORE POINT and SEQUENCE THREAD

upvoted 3 times

samm 5 years, 7 months ago

BDEF all 4 points are necessary to perform point in time recovery. Tricky one :) B. DB should be in mount mode D. Must be archive log mode. E. we need to specify SCN or Time to perform incomplete recovery F. need to be open in resetlogs.

upvoted 1 times

tamagogo 5 years, 6 months ago

E is not necessary. you can recover until cancel

upvoted 1 times

tlcecertified 5 years, 1 month ago

E is wrong, You can recover to an SCN, time, log sequence number, or restore point.

upvoted 5 times

Question #45 *Topic 1*

Which three are true about the SQL Tuning Advisor? (Choose three.)

A. It checks each query being analyzed for stale statistics.

B. It checks each query being analyzed for missing statistics.

C. It only recommends syntactic changes to SQL statements.

D. It can recommend semantic changes to SQL statements.

E. It considers all SQL statements being analyzed by the advisor task as a group.

F. It builds SQL profiles for each poorly performing SQL statement to prevent regressions.

Suggested Answer: *A, B, F*

Reference: https://docs.oracle.com/database/121/TGSQL/tgsql_sqltune.htm#TGSQL540

Community vote distribution



ald85 Highly Voted 6 years, 2 months ago

I think ABD

upvoted 14 times

markwu01 6 years, 1 month ago

Agree with ABD

upvoted 3 times

rodriguesJC 4 years, 9 months ago

Semantic changes????? I think is ABF, because can build SQL profiles

upvoted 2 times

freemun05 4 years, 5 months ago

depends on who SQL Prof. is runing, it's only true for auto mode , in manual it is only recomendation.

upvoted 1 times

dancymonkey 2 years, 6 months ago

ABD

https://docs.oracle.com/en/database/oracle/oracle-database/19/tgsql/sql-tuning-advisor.html#GUID-CB843CA1-6160-4918-BC71-86E36775EF15 AB - If these statistics are stale or missing, then the optimizer can generate suboptimal plans. Automatic Tuning Optimizer checks for missing or stale statistics, and recommends gathering fresh statistics if needed.

D - During structural analysis, Automatic Tuning Optimizer tries to identify syntactic, semantic, or design problems that can lead to suboptimal performance. The goal is to identify poorly written SQL statements and to advise you how to restructure them.

not c - only recommends syntactic changes

upvoted 1 times

veginha Highly Voted 6 years, 2 months ago

i think this is well answered, ABF

upvoted 9 times

Yennism 2 years, 7 months ago

Selected Answer: ABF

i think this is well answered, ABF

upvoted 1 times

Kuraudio 2 years, 7 months ago

D should be correct according to: https://docs.oracle.com/en/database/oracle/oracle-database/19/tgsql/sql-tuning-advisor.html#GUID-982EB0FC-AADD-4741-B0AA-79005C3AF657

upvoted 1 times

guretto 2 years, 8 months ago

A B F: "F" because the question is related to "SQL Tuning Advisor" not to the "Automatic" one. Therefore the following applies: "SQL profiles are created when a DBA invokes SQL Tuning Advisor." (https://docs.oracle.com/en/database/oracle/oracle-database/19/tgsql/managing-sql-profiles.html#GUID-1AC3C45E-331A-4016-8FE8-7FEFFD8D7220)

"D" is false because it can recommend syntactic, not semantic, restructuring

upvoted 1 times

auwia Most Recent 2 years, 1 month ago

F is saying for each poor ... D. Is saying "can" Probably they are A, B, D the correct answers.

upvoted 1 times

gio 3 years, 3 months ago

Selected Answer: ABD

i think abd

upvoted 1 times

flaviogcmelo 3 years, 11 months ago

Selected Answer: ABD

https://docs.oracle.com/en/database/oracle/oracle-database/19/tgsql/sql-tuning-advisor.html#GUID-73AB38C1-A7F6-401E-9010-B4476E173673

upvoted 2 times

egore0496 4 years, 7 months ago

Not F "for improving performance of high-load SQL statements" NOT "for each poorly performing SQL statement"
https://docs.oracle.com/en/database/oracle/oracle-database/18/tgsql/sql-tuning-advisor.html#GUID-73AB38C1-A7F6-401E-9010-B4476E173673
SQL Tuning Advisor is a mechanism for resolving problems related to suboptimally performing SQL statements.

Use SQL Tuning Advisor to obtain recommendations for improving performance of high-load SQL statements, and prevent regressions by only executing optimal plans.

upvoted 3 times

rashedptdba 4 years, 10 months ago

Correct:A,B,D

upvoted 4 times

Neil107 5 years ago

It's ABF.

upvoted 1 times

Manojx 5 years, 2 months ago

ABD https://docs.oracle.com/en/database/oracle/oracle-database/18/tgsql/sql-tuning-advisor.html#GUID-73AB38C1-A7F6-401E-9010-B4476E173673

upvoted 6 times

PortlandFighters 4 years, 10 months ago

Agree, it is ABD.

upvoted 2 times

test202102 5 years, 6 months ago

I vote for ABD

upvoted 3 times

janw 6 years ago

abf it is!

upvoted 1 times

janw 5 years, 12 months ago

i change to ABD! it will not create a profile for EACH poorly, e.g. one time queries. further: µµThe Automatic Tuning Optimizer identifies common problems with structure of SQL statements than can lead to poor performance. These could be syntactic, semantic, or design problems with the statement.

upvoted 8 times

cerebro2000x 6 years ago

abf is correct

upvoted 3 times

marcinb32 6 years, 1 month ago

ABF https://docs.oracle.com/en/database/oracle/oracle-database/18/tgsql/sql-tuning-advisor.html#GUID-6677340C-7DB3-4F0E-9921-46687E5EDC78

upvoted 2 times

klever35 6 years, 1 month ago

B and D

upvoted 1 times

Question #46 *Topic 1*

Which two are true about duplicating pluggable databases (PDBs) with RMAN? (Choose two.)

A. Two or more PDBs can be duplicated with the same RMAN DUPLICATE command.

B. All tablespaces belonging to a PDB must be duplicated when duplicating the PDB.

C. The auxiliary instance is automatically created with ENABLE_PLUGGABLE_DATABASE = TRUE.

D. A user with SYSDBA or SYSBKUP must use RMAN logged in to a PDB in order to duplicate it.

E. CDB$ROOT and PDB$SEED are automatically duplicated if a PDB is duplicated.

Suggested Answer: *C, D*

Reference: https://docs.oracle.com/database/121/BRADV/rcmdupdb.htm#BRADV430

Community vote distribution

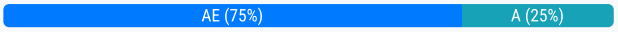

SimoneF Highly Voted 5 years, 3 months ago

A-E are correct:

"Duplicates the specified PDBs in the CDB. Use a comma-delimited list to duplicate multiple PDBs."

"When you duplicate one or more PDBs, RMAN also duplicates the root (CDB$ROOT) and the seed database (PDB$SEED)."

The reasons why the other answers are wrong are also found here: https://docs.oracle.com/database/121/BRADV/rcmdupdb.htm#BRADV760

upvoted 19 times

wiprooracle Highly Voted 5 years, 2 months ago

A & E https://docs.oracle.com/en/database/oracle/oracle-database/19/bradv/rman-duplicating-databases.html#GUID-68742310-3B60-4246-9431-6671697AB516

upvoted 7 times

mikimikimiki Most Recent 11 months, 3 weeks ago

Selected Answer: A

in question 174, people say that only one pdb could be cloned at a time, does it mean contradiction?

Question #174 Which three are true about cloning databases into an existing CDB with RMAN? (Choose three.)

A. A user with SYSDBA or SYSBKUP must be logged in with RMAN to the non-CDB being cloned when cloning a non-CDB. B. The target and auxiliary instance passwords must be the same. C. Two or more PDBs can be cloned with the same command. D. A non-CDB can be cloned into a PDB. E. Only one PDB can be cloned at a time. F. A user with SYSDBA or SYSBKUP must be logged in with RMAN to CDB$ROOT when cloning a PDB.

upvoted 1 times

auwia 2 years, 1 month ago

Selected Answer: AE

A. True B. False: "can" be done; and not "must" be done... C. False: it is started, and not it is created. Creation is manual. D. False: must login into CBD, and not PDB. E. True

upvoted 1 times

Yennism 2 years, 7 months ago

https://docs.oracle.com/en/database/oracle/oracle-database/19/bradv/rman-duplicating-databases.html#GUID-539E89F6-C0BC-49CB-8D8C-CD8FF88BAF32 According to last link: BC A: NO=> Only one PDB can be duplicated at a time. B: SI => Partial PDB duplication is not supported, only complete PDB duplication is supported. For example, you cannot include or exclude specific tablespaces while duplicating a PDB. C: SI => You must create the auxiliary instance as a CDB. To do so, start the instance with the following declaration in the initialization parameter file: enable_pluggable_database=TRUE D: NO => You must log in to the root of the CDB as a user who is granted the SYSDBA or SYSBACKUP privilege E:NO => Only If I duplicate a CDB

upvoted 1 times

Kuraudio 2 years, 7 months ago

You can select multiple PDBs to duplicate as single rman command according to this: https://docs.oracle.com/en/database/oracle/oracle-database/19/bradv/rman-duplicating-databases.html#GUID-571FD3C7-BF5A-4CF4-970A-9F8CAC3DF1F1

So A is true

upvoted 1 times

kaka321 3 years, 1 month ago

AE absolutely right https://docs.oracle.com/database/121/BRADV/rcmdupdb.htm#BRADV760 just above table25-1

upvoted 1 times

gio 3 years, 2 months ago

Selected Answer: AE

https://docs.oracle.com/en/database/oracle/oracle-database/19/rcmrf/DUPLICATE.html#GUID-E13D8A02-80F9-49A2-9C31-92DD3A795CE4 C is tricky: the auxiliary instance cannot be CREATED automatically. Auxiliary must be started in nomount with enable_pluggable_database=TRUE

upvoted 1 times

Nonicknameavailable 3 years, 5 months ago

why are some of you referencing 12.1 manuals for a 19c exam ??

upvoted 1 times

vkra 3 years, 6 months ago

Selected Answer: AE

A + E is correct, read this https://docs.oracle.com/database/121/BRADV/rcmdupdb.htm#BRADV760

upvoted 1 times

MartinDBA2 4 years, 2 months ago

Question is incomplete, it should state that the duplication is in a new CDB. If that is the case then the correct answers are A&E

upvoted 6 times

antjvr 4 years, 12 months ago

BE are correct: Duplicating a PDB to an existing CDB is subject to certain restrictions.

Only active database duplication is supported.

Only the following clauses of the DUPLICATE command are supported: NORESUME, DB_FILE_NAME_CONVERT, SECTION SIZE, and USING COMPRESSED BACKUPSET.

The following clauses of the DUPLICATE command are not supported: SPFILE, NO STANDBY, FARSYNC STANDBY, and LOG_FILE_NAME_CONVERT.

Duplicating a PDB to a CDB that is a standby database is not supported.

Only one PDB can be duplicated at a time.

Partial PDB duplication is not supported, only complete PDB duplication is supported. For example, you cannot include or exclude specific tablespaces while duplicating a PDB.

Duplicating a non-CDB as a PDB in an existing CDB is not supported.

Duplicating PDBs that contain TDE-encrypted tablespaces is not supported.

upvoted 2 times

marpayer 5 years ago

A & C https://docs.oracle.com/en/database/oracle/oracle-database/19/bradv/rman-duplicating-databases.html#GUID-8AEC7767-8572-435C-9320-8B589CA4F224 A. Two or more PDBs can be duplicated with the same RMAN DUPLICATE command. B. All tablespaces belonging to a PDB must be duplicated when duplicating the PDB. WRONG C. The auxiliary instance is automatically created with ENABLE_PLUGGABLE_DATABASE = TRUE. D. A user with SYSDBA or SYSBKUP must use RMAN logged in to a PDB in order to duplicate it. WRONG you must to connect to the CDB$ROOT not to the CDB E. CDB$ROOT and PDB$SEED are automatically duplicated if a PDB is duplicated. WRONG - if you duplicate into an existing PDB

upvoted 3 times

marpayer 5 years ago

sorry C is wrong you must start the auxiliary instance

upvoted 2 times

auwia 2 years, 1 month ago

Start is automatic (with that parameter set to true) , creation is manual!

upvoted 1 times

marpayer 5 years, 2 months ago

I think is E wrong, if you duplicate the database to an existing CDB, CDB$ROOT and PDB$SEED are not duplicated. And I think D is correct, you need SYSDBA or SYSBACKUP privileges to backup with RMAN https://docs.oracle.com/en/database/oracle/oracle-database/19/bradv/rman-duplicating-databases.html#GUID-F31F9FCE-B610-49EB-B9DB-44B9AA4E838F

upvoted 2 times

xRodge 5 years ago

From your own link, this is a quote that proves that E must be correct:

"When you duplicate one or more PDBs, RMAN also duplicates the root (CDB$ROOT) and the CDB seed (PDB$SEED). The resulting duplicate database is a fully functional CDB that contains the root, the CDB seed, and the duplicated PDBs."

upvoted 4 times

ObserverPL 5 years, 3 months ago

Agree - A and E are correct according to the official Oracle documentation.

upvoted 3 times

Question #47 *Topic 1*

Which three are true about Rapid Home Provisioning (RHP), which has been available since Oracle 18c? (Choose three.)

A. It is an Oracle Database service.

B. It can be used to deploy Oracle Database homes.

C. It cannot be used to provision applications.

D. It can be used to deploy Grid Infrastructure homes, including Oracle Restart.

E. It cannot be used to provision middleware.

F. It is an Oracle Grid Infrastructure service.

Suggested Answer: *D, E, F*

Reference: https://blogs.oracle.com/db_maintenance/rapid-home-provisioning-server-minimun-requirement

Community vote distribution

BDF (100%)

Aldrid Highly Voted 4 years, 9 months ago

Correct B D F

The architecture within a data center consists of an Oracle Grid Infrastructure (12.2 or later) deployment running the Rapid Home Provisioning centralized Service(also referred to as the RHP Server, but in fact merely a service of the Grid Infrastructure)

Rapid Home Provisioning and Maintenance enables you to install clusters, and provision, patch, scale, and upgrade Oracle Grid Infrastructure, Oracle Restart, and Oracle Database homes. The supported versions are 11.2, 12.1, 12.2, and 18c. You can also provision applications and middleware using Rapid Home Provisioning.

upvoted 14 times

Amaora 4 years, 7 months ago

about d :

Rapid Home Provisioning does not provision, patch, or upgrade Oracle Restart https://docs.oracle.com/en/database/oracle/oracle-database/12.2/cwadd/rapid-home-provisioning.html

upvoted 1 times

gio 4 years, 2 months ago

It can in 19c https://docs.oracle.com/en/database/oracle/oracle-database/19/cwadd/rapid-home-provisioning.html#GUID-492BAB58-26A7-4968-98BC-2993DDCCA856

upvoted 3 times

Yennism Most Recent 2 years, 1 month ago

Support for Single Instance Databases

You can use Rapid Home Provisioning to provision, patch, and upgrade single-instance databases running on clusters or Oracle Restart, or on single, standalone nodes. D is true. BDF https://docs.oracle.com/en/database/oracle/oracle-database/12.2/cwadd/rapid-home-provisioning.html#GUID-83518C02-E439-49E7-A406-47FE5DCF49B0

upvoted 1 times

ismoil 2 years, 4 months ago

B,D,E I THINK

upvoted 1 times

gio 2 years, 8 months ago

Selected Answer: BDF

i think

upvoted 1 times

dc_data 3 years, 1 month ago

Selected Answer: BDF

BDF [RHP 18c] Centralized repository of Gold Images – Grid Infrastructure, Database, application, middleware and generic homes

upvoted 2 times

eskias 3 years, 6 months ago

ans BDF A False ----- It is an Oracle Grid Infrastructure service B True -------- It can be used to deploy Oracle Grid Infrastructure and Oracle Database homes C False--------It can provision applications ans middle-ware D True----------RHP enables mass deployment and maintenance of standard operating environments of Oracle Database and Oracle Grid Infrastructure, including Oracle Restart. E False---------it can provision applications ans middle-ware F True----------It is an Oracle Grid Infrastructure service

upvoted 1 times

chito_s 4 years ago

RHP can provision applications and middleware. C says it can not. Similarly E says it can not. @arpz if E is wrong, C is also wrong. https://docs.oracle.com/en/database/oracle/oracle-database/12.2/cwadd/rapid-home-provisioning.html#GUID-CCEC5960-EDA4-4A3F-9643-0CA308EA49AA

upvoted 2 times

arpz 4 years, 5 months ago

Correct B C F

WRONG - A - RHP is NOT Oracle DB service

CORRECT - B - RHP can be used to deploy Oracle DB home and Grid Home

CORRECT - C - RHP can provision applications and middleware

WRONG - D - RHP is not supported for Oracle Restart

WRONG - E - RHP can provision applications and middleware

CORRECT - F - RHP is a service in Oracle Grid Infrastructure

Reference : https://docs.oracle.com/en/database/oracle/oracle-database/19/cwadd/rapid-home-provisioning.html#GUID-67917A2D-9C1E-423D-9BCB-FDD88E5663DD

upvoted 2 times

Question #48 *Topic 1*

Examine this configuration:

1. CDB1 is a container database.
2. COMMON_USER_PREFIX is C##.
3. PDB1 is a pluggable database contained in CDB1.
4. APP1_ROOT is an application container contained in CDB1.
5. APP1_PDB1 is an application PDB contained in APP1_ROOT.

You execute these commands successfully:

Which two are true? (Choose two.)

```
$ sqlplus sys/oracle_4U@localhost:1521/cdb1 as sysdba

SQL> CREATE USER c##user1 identified by oracle_4U container=all;
User created.

SQL> ALTER SESSION SET CONTAINER=pdb1;
Session altered.

SQL> CREATE USER p1_user1 identified by oracle_4U;
User Created.

SQL> ALTER SESSION SET CONTAINER=app1_root;
Session altered.

SQL> ALTER PLUGGABLE DATABASE APPLICATION app1_cdb1_app BEGIN INSTALL '1.0';
Session altered.

SQL> CREATE USER app1_user1 IDENTIFIED BY oracle_4U;
User Created.

SQL> ALTER PLUGGABLE DATABASE APPLICATION app1_cdb1_app END INSTALL '1.0';
Pluggable database altered.
```

A. APP1_USER1 can be created in PDB1.

B. APP1_USER1 can be created in CDB1.

C. APP1_USER1 can have different privileges in each Application PDB contained in APP1_ROOT.

D. C##_APP_USER1 can be created in CDB1.

E. P1_USER1 can be created in CDB1.

F. C##_USER1 will have the same privileges and roles granted in all PDBs in CDB1.

Suggested Answer: *C, F*

Official Answer: CF

Community vote distribution

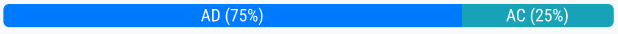

marcinb32 Highly Voted 6 years, 1 month ago

for me AD is correct.

upvoted 14 times

RinD Highly Voted 5 years, 4 months ago

A and D

For the common user: The username must be unique across all containers. So D is correct because APP_USER1 is different from APP1_USER1 in APP1_ROOT

B,E are wrong because you can only create local user in specific PDB.

A is true because APP1_USER1 is unique in PDB1

F is false because common user will not have the same privileges and role unless you grant it with option CONTAINER=ALL;

upvoted 5 times

RinD 5 years, 4 months ago

C is wrong because you can only grant privileges in application container and sync to all pdbs belong to it

upvoted 4 times

df404e2 Most Recent 1 year, 9 months ago

C can be done, as you technically are able to sync a different version of the application to every application PDB.

upvoted 1 times

Guhborges 3 years, 1 month ago

Selected Answer: AC

A,C. The letter D makes no sense, c##_app1 can be created at ANY TIME in cdb1, even without installing the patch in the approt

upvoted 1 times

gio 3 years, 2 months ago

Selected Answer: AD

ACD are all correct

upvoted 2 times

gio 3 years ago

A just in case PDB is not synchronized with app cdb

upvoted 1 times

vkra 3 years, 6 months ago

Selected Answer: AD

for me A+D is correct.

upvoted 1 times

flaviogcmelo 4 years, 9 months ago

A. APP1_USER1 can be created in PDB1. - Correct. It can be created - Tested B. APP1_USER1 can be created in CDB1. - Wrong. It must be a common user to be created in the CDB1 C. APP1_USER1 can have different privileges in each Application PDB contained in APP1_ROOT. - Correct - You can grant differen privilege to a local user on application pdbs D. C##_APP_USER1 can be created in CDB1. - Correct. It has the common user prefix, but make no sense. E. P1_USER1 can be created in CDB1. - Wrong. It must be a common user to be created in the CDB1 F. C##_USER1 will have the same privileges and roles granted in all PDBs in CDB1. - Wrong. It may have different privileges in PDBs in CDB1

In my opinion there are three correct answers in this question: A C D

upvoted 3 times

rodriguesJC 4 years, 9 months ago

Answer: AC D and F: C##_ not C## B and E: common users in CDB

upvoted 2 times

freemun05 4 years, 4 months ago

SQL> create user C##_APP_USER1 identified by q12345;

User created.

upvoted 1 times

wqui 4 years, 10 months ago

I think C is correct, becouse when you sync the PDB application, the user is created and we can grant whatever privilege we want, even in the next sync the privileges remaing forever until explicit revoke.

upvoted 1 times

dasigrist 5 years, 4 months ago

Is F wrong because the user created was C##USER1 and not C##_USER1?

upvoted 1 times

SimoneF 5 years, 3 months ago

Also, you can grant, for example, a privilege to a common user on the CDB$ROOT without specifying CONTAINER=ALL, meaning that privilege won't be replicated on the PDBS

upvoted 1 times

asefa 5 years, 6 months ago

C and F A ---PDB1 is pluggable db but APP1_USER1 is being created in App1_root not pdb B-- again the user is being created in application database not CDB

upvoted 3 times

janw 6 years ago

very strange question for me, do they mean 'has been created' instead of can... B,E for sure incorrect due to c## prefix

upvoted 1 times

janw 5 years, 12 months ago

checked: A, C, D all ok. but only 2 are requested.... i would say C incorrect as that user exists after the app installation, which is not done in the example script

upvoted 5 times

Alejandrrro 5 years, 6 months ago

I also wasn't sure about D and agree with you that it may be incorrect due to the fact that a PDB wasn't synced after the app installation. So the user wasn't synced to a pdb.

upvoted 1 times

cerebro2000x 6 years ago

DF for me

upvoted 1 times

julica 6 years ago

C and F looks ok. A,B - APP1_USER1 is a local user and from output is created in app1_cdb1_app and not in PDB1 or CDB1 D - C##_APP_USER1 there is nothing about such user in output E - P1_USER1 - based on the output, is a user created in PDB1.

If try from CDB: SQL> create user app1_user2 identified by oracle_4U; create user app1_user2 identified by oracle_4U * ERROR at line 1: ORA-65096: invalid common user or role name

SQL>

upvoted 4 times

dacoben415lywenw 6 years, 1 month ago

ACD all should work. E is not true, because of the prefix issue others also are not true, which is checked on 19c

upvoted 2 times

you1234 6 years, 1 month ago

its possible to create normal user in CDB? i think always we need create user like in CDB user C## prefix

upvoted 1 times

klever35 6 years, 1 month ago

for me E F are correct

upvoted 1 times

Question #49 *Topic 1*

Which two are facets of performance planning that should always be considered or implemented for an Oracle Database environment? (Choose two.)

A. defining primary keys for all tables to speed up all queries

B. using check constraints to speed up updates

C. defining foreign keys for all tables to speed up joins

D. the physical data model

E. the configuration of storage arrays

Suggested Answer: *A, E*

Official Answer: AE

Community vote distribution

DE (100%)

erial Highly Voted 5 years, 5 months ago

a,b and c are options to enforce data integrity not to speed up queries, i think the correct answers are d and e. if we have a good physical data model in an good storage array configuration we can think of good performance.

upvoted 12 times

Mastino69 5 years, 1 month ago

...and all constraints have a cost in terms of performance (index update and delete, checks included, etc.) to be managed by the instance itself.

upvoted 2 times

pgtips Highly Voted 5 years, 6 months ago

A. defining primary keys for all tables to speed up all queries B. using check constraints to speed up updates C. defining foreign keys for all tables to speed up joins D. the physical data model E. the configuration of storage arrays >> I believe D & E.

D - is to do with design the tablespace, partitions etc.. E - is to do with storage layout, that RAID levels e.g. 5 to keep data with high read IO. 1+0 for better write.

upvoted 8 times

pgtips 5 years, 6 months ago

A, if you think about warehouse table, not all tables are going to have primary key column (to save the storage).

upvoted 1 times

Yennism Most Recent 2 years, 1 month ago

Selected Answer: DE

correct answers are d and e

upvoted 1 times

ScottL 2 years, 9 months ago

Selected Answer: DE

DE correct

upvoted 1 times

flaviogcmelo 3 years, 9 months ago

Selected Answer: DE

https://docs.oracle.com/en/database/oracle/oracle-database/19/tgdba/IO-configuration-and-design.html#GUID-96CEF863-67AC-47BE-8834-8AB42864FC6E A good physical model will consider A, B and C options.

upvoted 2 times

brolem 3 years, 10 months ago

D, E. As stated, not all tables will have a, b, and c. Per https://docs.oracle.com/en/database/oracle/oracle-database/19/tgdba/IO-configuration-and-design.html#GUID-96CEF863-67AC-47BE-8834-8AB42864FC6E: Every Oracle database reads or writes data on disk, thus generating disk I/O. The performance of many software applications is inherently limited by disk I/O. Applications that spend majority of their CPU time waiting for I/O activity to complete are said to be I/O-bound. Additionally, a solid physical model will map out where and what is needed: The Physical Data Model of the Oracle Communications Data Model is the physical manifestation of the logical data model into database tables and relationships (or foreign key

constraints). Partitions, indexes, and Materialized Views have been added to aid performance.
(https://docs.oracle.com/cd/E84720_01/CDMRF/physical-data-model.htm#CDMRF717)
upvoted 1 times

mtnetmaker 3 years, 10 months ago
What is Oracle data model? The Physical Data Model of the Oracle Communications Data Model is the physical manifestation of the logical data model into database tables and relationships (or foreign key constraints). Partitions, indexes, and Materialized Views have been added to aid performance. so : D and E
upvoted 2 times

casfdsaf 3 years, 10 months ago
Selected Answer: DE
DE i think
upvoted 1 times

serkan110 4 years, 5 months ago
Primary Keys won't spenden up all queries
upvoted 1 times

Neil107 4 years, 6 months ago
I think it's DE Not ABC, since it all of them are a part of D.
upvoted 2 times

SimoneF 4 years, 9 months ago
"However, it is possible to make a sensible start by building indexes that enforce primary key constraints and indexes on known access patterns, such as a person's name."

"I/O Subsystem

The I/O subsystem can vary between the hard disk on a client PC and high performance disk arrays. Disk arrays can perform thousands of I/Os each second and provide availability through redundancy in terms of multiple I/O paths and hot pluggable mirrored disks. "

So I think A-E are quite viable answers after all, especially because the question asks what facet you should always "consider", which doesn't mean you should necessarily implement them all blindly. From: https://docs.oracle.com/cd/B28359_01/server.111/b28274/design.htm#CJHCJIDB
upvoted 2 times

RinD 4 years, 11 months ago
A,D

Refer to the question: "should always be considered or implemented" so E is correct but not always considered it due to company budget
upvoted 1 times

cerebro2000x 5 years, 6 months ago
DE if physcal model is about ram, cpu, os blocks, hba, switch etc
upvoted 6 times

CyborgQ 5 years, 6 months ago
A,E is correct: C - defining foreign keys not include an automatic index creation, therefore the foreign key can't have an index D- physical data model is a graphical schemas of table and relation
upvoted 1 times

tamagogo 5 years, 6 months ago
A and D define primary key will create index which will speed up queries
upvoted 2 times

klever35 5 years, 8 months ago
for sure C D is correct
upvoted 1 times

julica 5 years, 7 months ago
What does it mean more exactly "the physical data model"? DB Block? Better sounds the answer with "the configuration of storage arrays " than "the physical data model " :)
upvoted 1 times

Question #50 *Topic 1*

Which three actions are performed by Database Upgrade Assistant (DBUA)? (Choose three.)

A. It recompiles all stored PL/SQL code by using utlrp.sql.

B. It empties the RECYCLE BIN.

C. It performs prerequisite checks to verify if the Oracle database is ready for upgrade.

D. It sets all user tablespaces to ɹ€read-onlyɹ€ before starting the upgrade.

E. It removes the AUDSYS schema and the AUDIT_ADMIN and AUDIT_VIEWER roles

F. It increases tablespace size, if required, to meet upgrade requirements.

Suggested Answer: *A, C, F*

Reference: https://oracle-base.com/articles/12c/upgrading-to-12c

Community vote distribution



Rogazan Highly Voted 5 years, 2 months ago

https://docs.oracle.com/en/database/oracle/oracle-database/19/upgrd/upgrading-oracle-database-upgrade-assistant-dbua.html#GUID-8BDDF859-1638-45AC-824D-78FC8154847C A. False. Only "...This option recompiles all INVALID PL/SQL modules after the upgrade is complete B. True. "...The following is a list of examples of DBUA checks, and of actions that DBUA performs on the database: Empty database recycle bin..." C. True. "... DBUA runs the Pre-Upgrade Information Tool as part of the prerequisite checks it performs before starting the upgrade..." D. False. Only true for option Partial Backup "...Select among the following RMAN backup options:... Create a New Partial Offline RMAN Backup with User Tablespace. If you select this option, then user tablespaces are placed into read-only mode during the upgrade...." E. False. These elements are not mentioned in the description of the DBUA process. F. True. DBUA only continues if the check phase verifies that it can increase the size of tablespaces "... uses the following logic to modify or create new required tablespaces..."

upvoted 15 times

VinayakChandra 2 years, 1 month ago

I think A is correct because it clearly mentions "by using utlrp.sql" implying that only all invalid code will be recompiled

upvoted 1 times

asefa Highly Voted 5 years, 6 months ago

B,C,D come on guys A is not correct because it says "ALL stored PL/SQL CODE" DBUA recompiles only all invalid objects not all stored all stored PL/SQL F is also not correct It increases tablespace size, if required, to meet upgrade requirements. increasing tablespace size means adding more datafiles DBUA dont do this "If there is not enough disk space to grow, then DBUA prompts you to create space by adding more data files. DBUA does not automatically add new data files, because DBUA cannot determine where to create the files.: so B, C and D

upvoted 13 times

mamadu 2 years, 5 months ago

F is true: If the tablespaces are auto-extensible and the MAXSIZE initialization parameter needs adjustment, then DBUA prompts you to for this adjustment, and adjusts the MAXSIZE parameter. D is also true, but it is not performed implicitly by the DBUA unless you check that option unlike the case of F which is performed implicitly.

upvoted 1 times

JasonQian Most Recent 1 year, 5 months ago

D is incorrect, DBUA sets User Tablespaces to Read Only During the Upgrade, not before starting the upgrade https://docs.oracle.com/en/database/oracle/oracle-database/12.2/upgrd/upgrading-oracle-database-upgrade-assistant-dbua.html#GUID-29C7651A-28FD-4AA3-ACE4-D15708A9882B

upvoted 1 times

wagihov 1 year, 8 months ago

Selected Answer: BCF

according to documentation: A and E wrong B and C are correct. D is partially correct, as setting tablespaces to read only is an option that can be done during silent more. F is more correct as DBUA have a specific logic to modify or create new required tablespaces

https://docs.oracle.com/en/database/oracle/oracle-database/18/upgrd/upgrading-oracle-database-upgrade-assistant-dbua.html#GUID-8BDDF859-1638-45AC-824D-78FC8154847C

upvoted 1 times

zenai 1 year, 9 months ago

Selected Answer: BCF

guretto 9 months, 2 weeks ago F is true https://docs.oracle.com/en/database/oracle/oracle-database/19/upgrd/upgrading-oracle-database-upgrade-assistant-dbua.html#GUID-8BDDF859-1638-45AC-824D-78FC8154847C

upvoted 1 times

auwia 2 years, 1 month ago

Selected Answer: BCF

well documented.

upvoted 1 times

guretto 2 years, 7 months ago

F is true https://docs.oracle.com/en/database/oracle/oracle-database/19/upgrd/upgrading-oracle-database-upgrade-assistant-dbua.html#GUID-8BDDF859-1638-45AC-824D-78FC8154847C

DBUA uses the following logic to modify or create new required tablespaces:

If the data files are auto-extensible and have enough disk space to grow, then DBUA continues with the upgrade. If the data files are not auto-extensible, then DBUA prompts you and makes the files auto-extensible. If the tablespaces are auto-extensible and the MAXSIZE initialization parameter needs adjustment, then DBUA prompts you to for this adjustment, and adjusts the MAXSIZE parameter. If there is not enough disk space to grow, then DBUA prompts you to create space by adding more data files. DBUA does not automatically add new data files, because DBUA cannot determine where to create the files.

upvoted 1 times

Emperor11 2 years, 7 months ago

Selected Answer: BCD

I was searching about this one and i presume BCD are true, all in this link: https://docs.oracle.com/en/database/oracle/oracle-database/19/upgrd/upgrading-oracle-database-upgrade-assistant-dbua.html#GUID-F929338D-48EE-434F-A5F9-A9E7CD61B17F

B and C: The Prerequisite Checks window opens. DBUA analyzes the databases, performing preupgrade checks and displaying warnings as necessary. The following is a list of examples of DBUA checks, and of actions that DBUA performs on the database:

Empty database recycle bin.

Identify invalid objects.

Identify deprecated and desupported initialization parameters.

Identify time zone data file version.

D: Create a New Partial Offline RMAN Backup with User Tablespace. If you select this option, then user tablespaces are placed into read-only mode during the upgrade.

upvoted 1 times

musafir 2 years, 11 months ago

Selected Answer: BCD

BCD Are true all are B and C are definitely true but D is optional which is through RMAN. https://docs.oracle.com/en/database/oracle/oracle-database/19/upgrd/upgrading-oracle-database-upgrade-assistant-dbua.html#GUID-307DACD9-ECEE-4079-B767-B22620B99900

The thing about E it's mentioned that "If there is not enough disk space to grow, then DBUA prompts you to create space by adding more data files. DBUA does not automatically add new data files, because DBUA cannot determine where to create the files." It does not create the tablespace automatically but prompts to create the space.

upvoted 2 times

gio 3 years, 2 months ago

Selected Answer: BCF

i agree with Rogazan

upvoted 1 times

vkra 3 years, 6 months ago

Selected Answer: BCF

see comment from Rogazan and read the Oracle docs ...

upvoted 1 times

G_C 3 years, 10 months ago

Selected Answer: BCD

Question is not clearly formed A - FALSE I agree, DBUA could compile invalid objects (and not only PL/SQL, because synonyms too) B - TRUE - as a task of prerequisite mandatory checking task C - TRUE - this step is mandatory it perform always, but result can be skippable (using -ignorePreReqs option) D - TRUE - but ... it's OPTIONAL because we can choose f.g. "I have my own backup strategy" then D is FALSE E - FALSE F - FALSE it's only check but it will have never changed anything

upvoted 1 times

Heat418 4 years, 1 month ago

I vote for ABC: https://docs.oracle.com/en/database/oracle/oracle-database/12.2/upgrd/upgrading-oracle-database-upgrade-assistant-dbua.html#GUID-307DACD9-ECEE-4079-B767-B22620B99900

Using DBUA to Upgrade the Database on Linux, UNIX, and Windows Systems -> point 5

upvoted 1 times

Neil107 5 years ago

It is indeed ACF.

upvoted 2 times

Manojx 5 years, 2 months ago

ACF seems to be true. Language has some nuances but A is correct. B is wrong. It doesn't empty recycle bin by itself.

It increases tablespace is needed, but on prompt, not automatically. but it does increased the size.

upvoted 3 times

ssseee 5 years, 4 months ago

utlrp.sql compiles only invalid objects. So is not true

https://docs.oracle.com/en/database/oracle/oracle-database/18/rilin/recompiling-all-invalid-objects.html#GUID-1023C5FB-B409-4F16-BCA1-C14CB9B7E625

upvoted 2 times

Jatindra 5 years, 5 months ago

Yes correct answer is BCD ...

upvoted 2 times

dasigrist 5 years, 4 months ago

Hi, D has the same problem as F. It put the tablespace read-only "IF" you check the option. it's no automatic.

upvoted 1 times

SimoneF 5 years, 3 months ago

F has another problem though: DBUA doesn't directly increase the size of tablespaces, but rather it can enable autoextend or set maxsize to a higher value on their datafiles. In extreme cases: "If there is not enough disk space to grow, then DBUA prompts you to create space by adding more data files. DBUA does not automatically add new data files, because DBUA cannot determine where to create the files."

So in no case it directly resize tablespaces, hence I guess that D can be true and F false.

https://docs.oracle.com/en/database/oracle/oracle-database/12.2/upgrd/upgrading-oracle-database-upgrade-assistant-dbua.html#GUID-8BDDF859-1638-45AC-824D-78FC8154847C

upvoted 2 times

pgtips 6 years ago

A. It recompiles all stored PL/SQL code by using utlrp.sql. B. It empties the RECYCLE BIN. C. It performs prerequisite checks to verify if the Oracle database is ready for upgrade. D. It sets all user tablespaces to "read-only" before starting the upgrade. E. It removes the AUDSYS schema and the AUDIT_ADMIN and AUDIT_VIEWER roles F. It increases tablespace size, if required, to meet upgrade requirements.

I think it is ABC. F does not seem right, *DBUA will increase the size of TS _only_ if the data files are auto-extensible and have enough disk space to grow, then DBUA continues with the upgrade. If the data files are not auto-extensible, then DBUA prompts **you** and makes the files auto-

extensible. It is not capable to do by its own when auto-ext is disabled. Even if there are no space on the filesystem then DBUA prompts **you** to increase it.

upvoted 5 times

Question #51 *Topic 1*

Which two are true about RMAN backups when using a media manager to write backups to tape when there are only two tape drives? (Choose two.)

A. SBT tape compression can be used even if no RMAN compression is configured.

B. Any backup set written to the SBT device in this configuration can contain a maximum of two backup pieces.

C. Any backup written to the SBT device in this configuration can contain a maximum of two backup sets.

D. SBT tape compression and RMAN backup compression should be used in parallel.

E. The SBT device should be configured to use PARALLELISM 2 to allow both tape drive to be used simultaneously.

Suggested Answer: *D, E*

Official Answer: DE

Community vote distribution

AE (100%)

CyborgQ Highly Voted 4 years, 6 months ago

I think AE: B;C is false D: should be used in parallel...i can compress or Backup with RMAN> configure device type sbt backup type to compressed backupset; or via media manager sw for only sbt compression or together.....NOT should be used in parallel

upvoted 20 times

pgtips Highly Voted 4 years, 6 months ago

A. SBT tape compression can be used even if no RMAN compression is configured. B. Any backup set written to the SBT device in this configuration can contain a maximum of two backup pieces. C. Any backup written to the SBT device in this configuration can contain a maximum of two backup sets. D. SBT tape compression and RMAN backup compression should be used in parallel. E. The SBT device should be configured to use PARALLELISM 2 to allow both tape drive to be used simultaneously.

-- I think it is AE, A - sbt tape compression independently can be used on the TAPE side. Oracle recommends You should not use both tape compression provided by the media manager _and_ binary compression provided by RMAN. If the media manager compression is efficient, then it is usually the better choice. Using RMAN-compressed backup sets can be an effective alternative to reduce bandwidth used to move uncompressed backup sets over a network to the media manager, if the CPU overhead required to compress the data in RMAN is acceptable.

upvoted 12 times

gio Most Recent 1 year, 8 months ago

Selected Answer: AE

i agree with ae

upvoted 1 times

ScottL 1 year, 9 months ago

Selected Answer: AE

AE correct. Tested E

upvoted 1 times

vkra 2 years ago

Selected Answer: AE

A + E should be correct

upvoted 1 times

jackymak 2 years, 2 months ago

For backup set and backup pieces: A backup set contains one or more binary files in an RMAN-specific format. Each of these files is known as a backup piece. A backup set can contain multiple datafiles. For example, you can back up 10 datafiles into a single backup set consisting of a single backup piece. In this case, RMAN creates one backup piece as output. The backup set contains only this backup piece.
https://docs.oracle.com/cd/E18283_01/backup.112/e10642/rcmcncpt.htm#sect2

upvoted 1 times

Neil107 3 years, 6 months ago

It's AE

upvoted 5 times

mik8440p 4 years, 7 months ago

I think not D - "You should not use both tape compression provided by the media manager and binary compression provided by RMAN. If the media manager compression is efficient, then it is usually the better choice"
https://docs.oracle.com/cd/E11882_01/backup.112/e10642/rcmtunin.htm#BRADV90075

upvoted 3 times

klever35 4 years, 7 months ago

I think A D

upvoted 1 times

ald85 4 years, 8 months ago

I think BE

upvoted 1 times

ald85 4 years, 7 months ago

rectify... AE

upvoted 7 times

Question #52 *Topic 1*

You plan to install Oracle Grid Infrastructure for a Standalone Server and Oracle Database for the first time on a server.
Examine this command and its outcome:
Which two are true? (Choose two.)

```
# id oracle
uid=54321 (orcale) gid=54321(oinstall) groups=54321(oinstall), 54322 (dba)
```

A. oracle will be an owner of the Oracle Inventory.

B. oracle must be the owner of every Oracle Database installation.

C. oracle can own an Oracle Database installation but not an Oracle Grid Infrastructure installation.

D. oracle will be granted the SYSASM privilege when installing the Oracle Database software.

E. The user account, oracle, and group, oinstall, can be used for all Oracle software installations.

Suggested Answer: *C, D*
Official Answer: CD

Community vote distribution
AE (75%) CE (25%)

martinalexand Highly Voted 6 years, 1 month ago
A, E D not: The SYSASM Privilege for Administering Oracle ASM SYSASM is a system privilege that enables the separation of the SYSDBA database administration privilege from the Oracle ASM storage administration privilege. Access to the SYSASM privilege is granted by membership in an operating system group that is designated as the OSASM group. This is similar to SYSDBA and SYSOPER privileges, which are system privileges granted through membership in the groups designated as the OSDBA and OSOPER operating system groups. You can designate one group for all of these system privileges, or you can designate separate groups for each operating system privilege.
upvoted 19 times

imkissjr 3 years, 5 months ago
A is wrong - for Oracle Grid Infrastructure installation the owner of Oracle Inventory is the grid C , E are the right answers
upvoted 2 times

2211094 Most Recent 2 years, 2 months ago
But we ont have user grid on that question. Oracle can be used to install both oracle database and grid.
upvoted 1 times

rashedptdba Highly Voted 4 years, 10 months ago
Correct:A,E
upvoted 5 times

kaka321 3 years ago
you need to analyze this question based on the issued command. AE is correct.
upvoted 1 times

gio 3 years, 2 months ago
Selected Answer: CE
I think CE because if you install Oracle restart with grid user, grid user will be the owner of /u01/app/oraInventory
upvoted 1 times

gio 3 years, 2 months ago
Sorry AE, poorly written question but reasoning by exclusion AE
upvoted 1 times

ScottL 3 years, 3 months ago
Selected Answer: AE
AE correct
upvoted 1 times

vkra 3 years, 6 months ago
Selected Answer: AE

Reference: https://docs.oracle.com/en/database/oracle/oracle-database/19/cwlin/creating-operating-system-oracle-installation-user-accounts.html#GUID-4DA71FFA-AB5B-49AA-B21D-92C952C8A3FC

upvoted 1 times

jareach 3 years, 6 months ago

Selected Answer: AE

The others do not make sense.

upvoted 1 times

Neil107 5 years ago

I believe it's AC. E is a given thing, it goes back to 10g or even prior! D is incorrect.

upvoted 1 times

Rayyan0504 5 years ago

I will go with A,E

upvoted 2 times

asefa 5 years, 9 months ago

I will go with A, E

upvoted 4 times

logica 6 years, 1 month ago

for sure C is incorrect and E is correct

upvoted 1 times

ald85 6 years, 1 month ago

You are right, SYSASM is in Grid installation not in Database Software. AE are correct.

upvoted 5 times

ald85 6 years, 2 months ago

I think A,D

upvoted 1 times

veginha 6 years, 2 months ago

I was thinking in A, C

upvoted 1 times

ald85 6 years, 2 months ago

c is incorrect -> oracle can own database and grid in an single-user installation, and this output oracle has dba and oinstall group.

upvoted 3 times

proact 5 years, 10 months ago

C is correct as it saids, ".... but not Grid...."

upvoted 2 times

Question #53 *Topic 1*

Examine this configuration:

While CDB1 is open, `~/u02/app/oracle/fast_recover_area/cdb1/CDB1/controlfile02.ctl' is accidentally deleted.

To recover from this critical failure, you execute these commands:

What will be the outcome?

1. CDB1 is a container database running in ARCHIVELOG mode.
2. Controlfiles of CDB1 are multiplexed in '/u01/app/oracle/oradata/CDB1/controlfile/controlfile01.ctl' and '/u02/app/oracle/fast_recover_area/cdb1/CDB1/controlfile02.ctl'.
3. The only backup of CDB1 was taken when CONTROLFILE AUTOBACK was OFF
4. SNAPSHOT CONTROLFILE NAME is '/u01/app/oracle/product/12.2.0.1/db_1/dbs/snapcf_cdb1.f'.

```
$ rman target sys/oracle_4U@localhost:1521/cdb1

RMAN> SHUTDOWN ABORT
…
Oracle instance shut down

RMAN> STARTUP NOMOUNT

RMAN RESTORE CONTROLFILE FROM
'/u01/app/oracle/oradata/CDB1/controlfile/controlfile01.ctl';
```

A. It will create '$ORACLE_HOME/dbs/cdb1/CDB1/controlfile02.ctl'

B. It will create '/u01/app/oralce/oradata/CDB1/controlfile/controlfile02.ctl'.

C. It will re-create '/u02/app/oracle/fast_recover_area/cdb1/CDB1/controlfile02.ctl'

D. It will create '/u01/app/oracle/product/12.2.0.1/db_1/dbs/snapcf_cdb1control02.ctl'.

E. It will fail because there is no autobackup of the controlfiles.

Suggested Answer: *C*

Official Answer: C

Community vote distribution

C (100%)

niabs Highly Voted 4 years, 7 months ago

Agree with C

upvoted 11 times

Aldrid Highly Voted 4 years, 9 months ago

C - Correct

upvoted 9 times

gio Most Recent 2 years, 8 months ago

Selected Answer: C

i agree with c

upvoted 1 times

hilaire 3 years, 2 months ago

C correct

upvoted 2 times

Hacket 4 years ago

Selected Answer: C

C- Correct

upvoted 3 times

Question #54 *Topic 1*

Examine these actions:
1. Create a new database for a recovery catalog.
2. Create a tablespace with sufficient space in the catalog database for the recovery catalog.
3. Configure ARCHIVELOG mode for the catalog database.
4. Create a user to own the recovery catalog schema with quota on the tablespace that will contain the catalog.
5. Grant the RECOVERY_CATALOG_OWNER role to the recovery catalog schema owner.
6. Grant the SYSBACKUP privilege to the recovery catalog schema owner.
Which are the minimum actions that must be performed before executing the CREATE CATALOG command?

A. 2, 4, 5, 6

B. 1, 2, 3, 4, 5, 6

C. 1, 2, 4, 5

D. 2, 4, 5

E. 1, 3, 4, 5

Suggested Answer: *D*
Official Answer: D

Apollo5 Highly Voted 6 years ago
D is the right answer. Step No 2 is required, so it must be contained in the choise Step 1 is optional Step 3 is not necessary Steps 4 & 5 are required Step 6 is not necessary
upvoted 18 times

Guhborges 3 years, 1 month ago
Selected Answer: D
2-4-5 letter D is Right.

1 is wrong, cause you don't necessary need a new database to be the catalog. :)
upvoted 1 times

gio 3 years, 2 months ago
Selected Answer: D
I agree with Apollo5. Recovery catalog is a schema that can be contained in any existing DB
upvoted 1 times

ScottL 3 years, 3 months ago
Selected Answer: D
D - see Apollo5
upvoted 1 times

vkra 3 years, 6 months ago
Selected Answer: C
I would say C too
upvoted 1 times

Nonicknameavailable 3 years, 6 months ago
Correct answer is C. The question gives a sequence of steps and asks you to examine the minimum options in that sequence. It does not tell you a database already exists.
upvoted 2 times

auwia Most Recent 2 years, 1 month ago
The problem is saying "create a new database", to be correct your sentence it should be "create a database" --> not new
upvoted 2 times

senator 4 years, 9 months ago
D is the correct answer when reading both 12c and 19c oracle docs. https://docs.oracle.com/en/database/oracle/oracle-database/19/rcmrf/CREATE-CATALOG.html#GUID-9CFE47C7-2273-48C4-8056-6A4042AF3CDC Examples

Example 2-70 Creating a Recovery Catalog and Registering a Database

Assume that you start SQL*Plus and connect to the recovery catalog catdb with administrator privileges. You execute the CREATE USER statement as follows, replacing password with a user-specified password (see Oracle Database Security Guide for information on creating secure passwords). The SQL statement creates a user rco in database catdb and grant the rco user the RECOVERY_CATALOG_OWNER role.

SQL> CREATE USER rco IDENTIFIED BY password 2 DEFAULT TABLESPACE cattbs 3 QUOTA UNLIMITED ON cattbs; SQL> GRANT recovery_catalog_owner TO rco; SQL> EXIT

upvoted 2 times

smartvan 4 years, 9 months ago

This is all you need to create the catalog.

tested

SQL> CREATE USER restoretests IDENTIFIED BY oracle >>>>>> create user(4) DEFAULT TABLESPACE tools >>>>>>>>>> create tablespace (2) TEMPORARY TABLESPACE temp QUOTA UNLIMITED ON tools;

SQL> GRANT CONNECT, RECOVERY_CATALOG_OWNER TO ; >>> grant privilege (5)

$ rman target / catalog / @ (instance name) select instance_name from v$instance;

RMAN> create catalog;

correct answer is D

upvoted 2 times

niabs 5 years, 1 month ago

I will choose C Because likely you will create a dedicated database for the catalog unless you already have a separated one, which is not mentioned here

upvoted 2 times

freemun05 4 years, 5 months ago

question is specific - "minimum actions" D is correct

upvoted 1 times

Rogazan 5 years, 2 months ago

I agree with option D (minimum of actions), but if Oracle asks me in the exam, I will answer that the correct option is C. This confirms that we know that the Catalog Database must be a different one which we want to protect!

upvoted 4 times

Jatindra 5 years, 4 months ago

Correct Ans: D

upvoted 2 times

asefa 5 years, 6 months ago

D is the minimum actions that must be performed before executing create catalog command

upvoted 4 times

pgtips 5 years, 11 months ago

I believe, E, as the Option 1: Create a new Database is one of the Choice I go for E. If it was not there then probably D makes more sense.

upvoted 2 times

SimoneF 5 years, 3 months ago

1 is disputable but I assume they imply you do have a DB, and you should create another one. In that case, it's not necessary to simply give the command, even though you wouldn't create a catalog on a database which is not dedicated. Besides, 3 - Enabling archivelog mode on the catalog database - is not necessary to create a catalog.

upvoted 1 times

janw 6 years ago

agree with d, only 4&5 really required.

upvoted 3 times

julica 6 years ago

Oly D is correct am minimal steps.

upvoted 2 times

klever35 6 years, 1 month ago

Correct Answer E

upvoted 1 times

Question #55 *Topic 1*

Which two are true about creating pluggable databases (PDBs) using snapshots in Oracle 19c and later releases? (Choose two.)

A. A PDB snapshot is always a full copy of the source PDB.

B. A PDB snapshot is always a sparse copy of the source PDB.

C. A snapshot copy PDB depends on a storage snapshot which can only be stored on specific file systems.

D. A PDB snapshot depends on a storage snapshot which can be stored on any file system.

E. A PDB snapshot depends on a storage snapshot which can only be stored on specific file systems.

F. A snapshot copy PDB depends on a storage snapshot which can be stored on any file system.

G. A snapshot copy PDB can be created from a stand-alone clone PDB.

Suggested Answer: *A, E*

Reference: https://docs.oracle.com/en/database/oracle/oracle-database/18/multi/administering-pdb-snapshots.html#GUID-35B03C37-00BA-4945-BF20- 9A7C4C541955

Community vote distribution

CG (100%)

ald85 Highly Voted 4 years, 8 months ago

I think CG http://ora-srv.wlv.ac.uk/oracle19c_doc/multi/introduction-to-the-multitenant-architecture.html#GUID-47847140-8DE4-4C4E-8C93-4E9C197D1402 A, B incorrects -> http://ora-srv.wlv.ac.uk/oracle19c_doc/multi/introduction-to-the-multitenant-architecture.html#GUID-47847140-8DE4-4C4E-8C93-4E9C197D1402 C correct, D,E,F incorrect -> https://docs.oracle.com/en/database/oracle/oracle-database/19/multi/cloning-a-pdb.html#GUID-B00A0E48-C892-4DC8-8455-C6F2ABC8EF91 G -> correct -> A clone from a PDB snapshot is a full, standalone PDB. -> https://docs.oracle.com/en/database/oracle/oracle-database/19/multi/cloning-a-pdb.html#GUID-22118625-3157-4C38-8A88-79A561A3E6ED

upvoted 19 times

gio Most Recent 1 year, 8 months ago

Selected Answer: CG

I'm sure of CG after looking at https://docs.oracle.com/en/database/oracle/oracle-database/18/multi/cloning-a-pdb.html#GUID-F26C8A55-2465-475A-905A-83BD2E68AE6E

upvoted 1 times

hilaire 2 years, 1 month ago

Pdb Snapshot dépends of Database storage CG

upvoted 2 times

Bio12345 2 years, 11 months ago

Selected Answer: CG

C G as argumented by others

upvoted 2 times

CyborgQ 4 years, 6 months ago

Correct is AB

upvoted 1 times

julica 4 years, 6 months ago

First of all they are talking about 19c+ versions. In 19c docs they say: A PDB snapshot is a point-in-time copy of a PDB. The source PDB can be open read-only or read/write while the snapshot is created. You can create snapshots manually using the SNAPSHOT clause of CREATE PLUGGABLE DATABASE (or ALTER PLUGGABLE DATABASE), or automatically using the EVERY interval clause. If the storage system supports sparse clones, then the preceding command creates a sparse copy. Otherwise, the command creates a full copy. So, A and B are false from my point of view.

upvoted 2 times

janw Highly Voted 4 years, 5 months ago

indeed A, B are false as you argumented well

in my opion, correct answer C, G

G: https://docs.oracle.com/en/database/oracle/oracle-database/18/multi/cloning-a-pdb.html#GUID-F26C8A55-2465-475A-905A-83BD2E68AE6E and then create a SNAPSHOT COPY PDB from the standalone PDB.

C: Storage requirements depend on the setting of the CLONEDB initialization parameter:

CLONEDB=FALSE

The underlying file system for the source PDB files must support storage snapshots. Such file systems include Oracle Automatic Storage Management Cluster File System (Oracle ACFS) and Direct NFS Client storage. Oracle Exadata supports snapshot copy functionality on ASM configured with sparse ASM grid disks.

CLONEDB=TRUE

The underlying file system for the source PDB files can be any local file system, network file system (NFS), or clustered file system that has Direct NFS enabled and supports sparse files. Most UNIX systems meet these requirements, including Oracle ACFS and ZFS. However, when CLONEDB=TRUE, the source PDB must remain in open read-only mode when clones exist.

upvoted 5 times

janw 4 years, 5 months ago

D, E false as a pdb snapshot is not dependent of the underlying storage snapshot technology

upvoted 1 times

julica 4 years, 6 months ago

Full and Sparse Snapshots

The content of snapshots generated by ALTER PLUGGABLE DATABASE ... SNAPSHOT depends on the underlying file system. If the underlying file system supports sparse copies, then the PDB-level snapshots are sparse. Only the first PDB-managed PDB snapshot is full. Otherwise, the PDB snapshots contain full copies of the data files. The snapshot includes other files necessary to create a PDB from the snapshot.

If the file system supports sparse files, then all PDB snapshots in the carousel except the first one are sparse. The source PDB can remain in read/write mode. Sparse files significantly reduce the carousel storage space.

upvoted 1 times

klever35 4 years, 8 months ago

I think B, E

upvoted 1 times

taotsumiau 4 years, 8 months ago

I think A, E

upvoted 1 times

veginha 4 years, 8 months ago

C, E? Anyone?

upvoted 1 times

Question #56 *Topic 1*

Which two are true about Oracle Database Configuration Assistant (DBCA) templates? (Choose two.)

A. The General Purpose of Transaction Processing templates are most suitable when concurrency and recoverability are key criteria.

B. Oracle DBCA templates can store only logical structure and not database files.

C. New templates can only be created by modifying an existing user-created template.

D. The Data Warehouse template is most suitable when transaction response time is the key criterion.

E. Oracle DBCA templates can be used to create new databases and duplicate existing databases.

Suggested Answer: *A, E*

Reference: https://docs.oracle.com/database/121/ADMQS/GUID-2B8A3B5E-D319-4377-8B22-8BB67DCC9885.htm#ADMQS0235

Community vote distribution

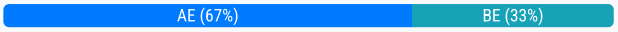

taotsumiau Highly Voted 6 years, 1 month ago

A&E are correct. The new features are in Oracle 19c https://docs.oracle.com/en/database/oracle/oracle-database/19/admqs/installing-oracle-database-creating-database.html#GUID-418CDFD6-340C-40B4-AC08-5005560ED9DC

upvoted 19 times

janw Highly Voted 6 years ago

agree with AE

upvoted 7 times

wagihov Most Recent 1 year, 8 months ago

Selected Answer: AE

Obviously A and E are correct, while C and D are wrong, and those four are already discussed. but what about B? according to documentation, "seed" templates can include datafiles, and that's why B is wrong. so the Correct answer is A and E https://docs.oracle.com/database/121/ADMQS/GUID-418CDFD6-340C-40B4-AC08-5005560ED9DC.htm

upvoted 1 times

gonza150 1 year, 9 months ago

Yes, Oracle Database Configuration Assistant (DBCA) templates can store database files, but it depends on the type of template you create:

Data Templates (with data files): These templates include both the database structure and actual database files, allowing you to replicate a fully functioning database with pre-configured tables, indexes, and any other objects that were in the original database. This type of template can be very large because it contains the physical data files.

Structure-Only Templates: These templates only contain the structural definitions of the database (such as tablespaces, parameters, and configuration) without the actual data files. When you create a database with a structure-only template, you're effectively creating a new database that mimics the structure of the original but does not include any data.

When you create or use a template in DBCA, make sure to specify whether you want to include database files or just the structure, depending on your needs.

upvoted 1 times

auwia 2 years, 1 month ago

Selected Answer: BE

A. False -> "The General Purpose of Transaction" it's "or" not "of" -> definition is definition, should be copy and paste :-) B. True -> Oracle DBCA templates are XML file, they cannot store data files! C. False D. False E. True

upvoted 1 times

vkra 3 years, 6 months ago

Selected Answer: AE

A + E is correct

upvoted 1 times

Dhanushbh 5 years ago

In Oracle 18c the DBCA can now duplicate container databases. In the following example we have a container database called "cdb1". We use the following command to attempt to create a new CDB called "cdb2" as a clone of "cdb1", but it fails because the source database is not in archivelog mode.

E correct
upvoted 2 times

Asmodeus 5 years, 4 months ago
Very tricky question. A: true. https://docs.oracle.com/en/database/oracle/oracle-database/19/admqs/installing-oracle-database-creating-database.html#GUID-E59E0AE9-AF91-4DD2-B311-C92E3EFE9948 General Purpose or Transaction processing Many concurrent users perform numerous transactions that require rapid access to data. Availability, speed, concurrency, and recoverability are key issues. B: false. Some templates store database files. C: false. I can create template from existing database too. D: true From link provided above: Data warehouse Users perform numerous, complex queries that process large volumes of data. Response time, accuracy, and availability are key issues. E: true https://docs.oracle.com/en/database/oracle/oracle-database/19/admqs/installing-oracle-database-creating-database.html#GUID-2B8A3B5E-D319-4377-8B22-8BB67DCC9885 Templates are used in DBCA to create new databases and duplicate existing databases.

So we have A and D on the same level (both are talking about existing templates) and E that point on DBCA usage at all, not the templates.

Sometimes I think that A&D is correct answer...
upvoted 3 times

Eric_F 4 years, 3 months ago
Option D: when transaction response time is the key criterion. but in doc, it describe as Response time, nor transaction response time
upvoted 2 times

gio 3 years, 2 months ago
Yes tricky question or poorly written question. I think A D for the same reasons
upvoted 1 times

SimoneF 5 years, 3 months ago
D is false, complex queries are not the same as transaction processing:

"A data warehouse is a database designed to enable business intelligence activities: it exists to help users understand and enhance their organization's performance. It is designed for query and analysis rather than for transaction processing"

https://docs.oracle.com/en/database/oracle/oracle-database/19/dwhsg/introduction-data-warehouse-concepts.html#GUID-452FBA23-6976-4590-AA41-1369647AD14D
upvoted 4 times

auwia 2 years, 1 month ago
As you wrote "General Purpose or Transaction processing", in the theory it is written"or", in the answer "of" so I think A is wrong!
upvoted 1 times

cerebro2000x 6 years ago
AE for me
upvoted 2 times

CyborgQ 6 years ago
AE correct
upvoted 3 times

you1234 6 years, 1 month ago
Klever35,

have you attended the exam 1z0-083? Please suggest
upvoted 3 times

klever35 6 years, 1 month ago
I think B E
upvoted 2 times

Question #57 *Topic 1*

Your SALES_ROOT application container has two application PDBs.
The SALES_APP application has a common table, FIN.REVENUE, in the two PDBs.
Examine this query and its output:
Which two are true? (Choose two.)

```
SELECT containers_default, container_map, table_name
  FROM dba_tables WHERE owner='FIN';

CONTAINERS_DEFAULT    CONTAINER_MAP    CONTAINER_MAP_OBJECT   TABLE_NAME
------------------- --------------    --------------------   ------------
NO                    YES               NO                    REVENUE
NO                    NO                YES                   MAPTABLE
```

A. The CONTAINERS clause cannot be used in queries on the REVENUE table.

B. The REVENUE table must be a list-partitioned table.

C. The MAPTABLE tables defines a logical partition key on a commonly used column for the REVENUE table.

D. The MAPTABLE table is a metadata-linked table.

E. A container map exists for the REVENUE table, but is not enabled.

F. The REVENUE table partitions are not pruned across the PDBs automatically.

Suggested Answer: *C, E*
Official Answer: CE

michael_live Highly Voted 5 years, 10 months ago
C and F are correct. Tested in lab with exact same setting like in the output shown.

A. -> incorrect - The CONTAINERS clause can be used. B. -> incorrect - It can be also a range based. C. -> CORRECT - https://docs.oracle.com/en/database/oracle/oracle-database/19/multi/administering-application-containers-with-sql-plus.html#GUID-AAF93A02-7C70-4024-8758-E351C213543E D. -> incorrect - Regular table E. -> incorrect - shown in the output CONTAINER_MAP = YES F. -> CORRECT - because CONTAINERS_DEFAULT is not enabled for the revenue table.
upvoted 13 times

auwia 2 years, 1 month ago
Because of what I can read from your link: "A container map can define a logical partition key on a column for a common object. " It seems to me that C is false! And the full statement is true if you invert the 2 object names! :-) MAPTABLE has CONTAINER_MAP=NO REVENUE has CONTAINER_MAP=YES
upvoted 1 times

auwia Most Recent 2 years, 1 month ago
Selected Answer: DF
For C. see my comment in reply to "michael_live" For D. See decoben415... comment below. For F. see comment from michael_live
upvoted 1 times

gio 3 years, 2 months ago
Selected Answer: CF
CF but michael_live reason for F is incorrect.
upvoted 2 times

Bio12345 4 years, 5 months ago
Selected Answer: CF
C F, thanks michael_live
upvoted 2 times

erial 5 years, 11 months ago
based on this: container_map (yes) and containers_default (yes), when used together, allow, container_map to prune the partitions and therefore the application pdbs, based on the key that is passed to the query correct answers are: c,f

upvoted 1 times

Chansi 5 years, 11 months ago

sorry E is wrong.. The object is enabled

upvoted 1 times

Chansi 5 years, 11 months ago

C and E as the Map_table is not enabled . The container_map column shows "NO"

upvoted 1 times

dacoben415lywenw 6 years, 1 month ago

Seems D is also correct

SQL> select containers_default, container_map, container_map_object, table_name from DBA_TABLES where table_name='CONTAINERMAP';

CON CON CON --- --- --- TABLE_NAME ---------- NO NO YES CONTAINERMAP

SQL> select SHARING from dba_objects where OBJECT_NAME='CONTAINERMAP' and OBJECT_TYPE='TABLE';

SHARING ---------- METADATA LINK

upvoted 2 times

veginha 6 years, 1 month ago

C and E -Use a container map to partition the data in metadata-linked objects. Container maps partition data in application PDBs based on a commonly-used column. -The map object is the partitioned table. The names of the partitions in the map table match the names of the application PDBs in the application container. The metadata-linked object is not physically partitioned at the table level, but it can be queried using the partitioning strategy used by the container map. -In addition, you can enable the CONTAINERS_DEFAULT attribute for a table or view in an application root. When this attribute is enabled, the CONTAINERS clause is used for queries and DML statements on the database object by default, and the CONTAINERS clause does not need to be specified in the SQL statements. To enable the CONTAINERS_DEFAULT attribute for a table or view in an application root, run the an ALTER TABLE or ALTER VIEW statement with the ENABLE CONTAINERS_DEFAULT clause.

https://docs.oracle.com/en/database/oracle/oracle-database/19/multi/administering-application-containers-with-sql-plus.html#GUID-AAF93A02-7C70-4024-8758-E351C213543E

upvoted 2 times

veginha 6 years, 1 month ago

i mean C and F

upvoted 2 times

Question #58 *Topic 1*

Your container database, CDB1, has an application container, HR_ROOT, with an application PDB, HR_PDB1.
You have the required privilege to clone HR_PDB1 to container database CDB2, which does not contain HR_ROOT.
Which two are always true? (Choose two.)

A. CDB1 and CDB2 must be in shared undo mode.

B. A common user must exist in CDB2 with the CREATE PLUGGABLE DATABASE privilege.

C. All transactions in HR_PDB1 of CDB1 must commit before the cloning process starts.

D. Cloning HR_ROOT automatically clones HR_PDB1.

E. The HR_PDB1 clone created in CDB2 will be in mount state when cloning ends.

Suggested Answer: *B, D*
Official Answer: BD

Community vote distribution
BE (100%)

mqdiri Highly Voted 5 years, 4 months ago
BE

After you create the application PDB, it is in mounted mode, and its status is NEW https://docs.oracle.com/en/database/oracle/oracle-database/21/multi/creating-application-pdb.html#GUID-88E61123-F651-47D7-9977-B589268EBD29
upvoted 12 times

2211094 Most Recent 2 years, 1 month ago
CE C===To ensure data integrity E===According to Oracle documentation, after the cloning process is completed, the cloned PDB will be in the mount state. This allows the DBA to perform any additional configurations or checks before opening the PDB.
https://docs.oracle.com/en/database/oracle/oracle-database/19/multi/cloning-a-pdb.html#GUID-D4900E43-DF38-4629-9A1F-D1C11EC8EC8F
upvoted 1 times

ismoil 2 years, 10 months ago
C E I think
upvoted 2 times

gio 3 years, 2 months ago
Selected Answer: BE
Sure of BE https://docs.oracle.com/en/database/oracle/oracle-database/19/multi/cloning-a-pdb.html#GUID-D4900E43-DF38-4629-9A1F-D1C11EC8EC8F
upvoted 1 times

ErikJanssen 3 years, 4 months ago
BC I think. After looking through https://oracle-base.com/articles/19c/multitenant-dbca-pdb-remote-clone-19c E is not correct, the remote pdb is open read write after being cloned. C could be true; Source database needs to be readonly if remote CDB is not in archivelog mode or does not have local undo enabled.
upvoted 1 times

jareach 3 years, 6 months ago
D, E in my opinion.
upvoted 1 times

Neil107 5 years ago
It's B and E.
upvoted 2 times

cerebro2000x 6 years ago
BE for me

you need to clone the app root first and the pdb [separately]

CREATE PLUGGABLE DATABASE hr_root_dev AS APPLICATION CONTAINER FROM hr_root@link_hr_root;

upvoted 2 times

cerebro2000x 6 years ago

CREATE PLUGGABLE DATABASE operations_dev FROM operations@link_operations

upvoted 1 times

CyborgQ 6 years ago

B,D are correct E incorrect: the clone can be done in read/write mode and the copy is in read/write mode

upvoted 1 times

SimoneF Highly Voted 5 years, 3 months ago

I didn't find any support to this claim. On the contrary, everything leads to think that, even in open mode, your cloned PDB will be mounted after the creation:

"After you create the PDB, it is in mounted mode, and its status is NEW. "

https://docs.oracle.com/en/database/oracle/oracle-database/21/multi/cloning-a-pdb.html#GUID-D4900E43-DF38-4629-9A1F-D1C11EC8EC8F

upvoted 5 times

veginha 6 years, 1 month ago

BE, but I'm not sure. Anyone know this?

upvoted 1 times

Question #59 *Topic 1*

Which three are true about backup, restore, and recovery operations done without using Recovery Manager (RMAN)? (Choose three.)

A. Backing up a database in NOARCHIVELOG mode using O/S utilities requires that the database instance be started and the database be in the MOUNT state.

B. Backing up a database in ARCHIVELOG mode using O/S utilities requires that the database instance be started and the database be in MOUNT state.

C. An Oracle database can be restored from backup files copied using O/S utilities.

D. Oracle data file backups, copied using an O/S utility, can be added to the RMAN catalog as IMAGE COPIES.

E. Backing up a database in NOARCHIVELOG mode using O/S utilities requires that the database instance be shut down.

F. Oracle archive log backups, copied using an O/S utility, can be added to the RMAN catalog as a backup set.

G. Backing up a database in ARCHIVELOG mode using O/S utilities requires that the database instance be started and the database be in OPEN state.

Suggested Answer: *B, D, E*
Official Answer: BDE

Community vote distribution
CDE (75%) CDF (25%)

janw Highly Voted 5 years ago
i go for CDE
upvoted 22 times

julica Highly Voted 5 years, 1 month ago
A,B,G are not right because the question is telling us 'without RMAN'. To put it in mount/open or so, are necessary only in case of using RMAN. Otherwise, using OS Utils, DB must be closed or maximum in NOMOUNT state.
upvoted 12 times

Mastino69 4 years, 7 months ago
Also F should be wrong because, you cannot add backup set to rman catalog but only backup pieces and image copies
upvoted 3 times

Alejandrrro 4 years, 6 months ago
Agree

F is incorrect as according to the following note only archivelog, backuppiece, controlfilecopy or datafilecopy can be cataloged

https://docs.oracle.com/cd/B19306_01/backup.102/b14194/rcmsynta011.htm
upvoted 1 times

hkshado 4 years, 9 months ago
Generally, I agree with you. However, back to the Oracle 9 when RMAN is not available. We can backup the DB (with archive log enabled) which is opened in read/write mode using OS copy command. (I am sure about that because my company used this for Oracle 9 online backup). what you need to do is something like following: 1. Execute "ALTER TABLESPACE tb_name BEGIN BACKUP;" for each tablespace before the backup start. 2. Copy files 3. Execute "ALTER TABLESPACE tb_name END BACKUP;" for each tablespace after the backup end.
https://docs.oracle.com/cd/B10501_01/server.920/a96572/osbackups.htm
upvoted 2 times

Mastino69 4 years, 7 months ago
Agree but it is not mentioned in G. It should have said: ...the database OPEN, and a command ALTER DATABASE... BEGIN BACKUP executed. Then you can copy the files with OS utils.
upvoted 1 times

SimoneF 4 years, 3 months ago
But answere G states that it is REQUIRED to have the database Open, while what you said proves you simply CAN do that with the db open. Nothing prevents backing up the database via OS copy when it's shutdown.

upvoted 5 times

1451755 Most Recent 1 year, 4 months ago

Selected Answer: CDF

CDF E - false, database can be in the MOUNT state OR in READ ONLY mode. In these modes db. data is not changing so, consistent backup is possible for NOARCHIVELOG db. F - true, I did it a few times successfully, copy backups with OS command and user RMAN CATALOG START WITH ... Doc: https://docs.oracle.com/en/database/oracle/oracle-database/21/rcmrf/CATALOG.html "START WITH 'string_pattern'" = "Catalogs all valid backup sets, data file and control file copies, and archived redo log files whose name start with string_pattern. The string pattern can be an ASM disk group, Oracle-managed files directory, or part of a file name ..."

upvoted 1 times

piontk 1 year, 11 months ago

Selected Answer: CDE

CDE correct. Refs: https://docs.oracle.com/en/database/oracle/oracle-database/19/bradv/user-managed-flashback-dbpitr.html#GUID-D71F3FF3-84E5-4D81-9307-E2AD5CB51833 https://docs.oracle.com/en/database/oracle/oracle-database/19/bradv/rman-backup-concepts.html#GUID-C0BD0706-29EF-44C6-AF53-401522B4BDD0

upvoted 1 times

gio 2 years ago

Selected Answer: CDF

I think CDF

upvoted 1 times

ScottL 2 years, 3 months ago

Selected Answer: CDE

C,D,E Correct

upvoted 1 times

jareach 2 years, 6 months ago

Definitely CDE

upvoted 2 times

Bio12345 3 years, 5 months ago

Selected Answer: CDE

CDE, most voted

upvoted 4 times

smartvan 3 years, 9 months ago

Captain here

Backing up a database with OS utilites can be performed only when the database is shut down, so... A, B, G are incorrect

OS copies are not backupsets, so again F is incorrect.

It remains CDE that should be correct.

upvoted 1 times

Neil107 4 years ago

The question is: '...about backup, restore, and recovery operations done WITHOUT using Recovery Manager (RMAN)?' Think only about the use of OS utilities here. I would disregard DF since it speaks about 'adding' the OS taken backup to RMAN. If RMAN is not an option at all, I believe the correct answers are CEG

upvoted 1 times

Schmeili 4 years, 9 months ago

E: Backing up a database in NOARCHIVELOG mode using O/S utilities requires that the database instance be shut down. --> OS-Backup is possible when the DB is in nomount/mount state, too ...so i think E is not correct

i go with CDF... why f: you can recatalog any backup to rman catalog....

upvoted 5 times

brolem 3 years, 4 months ago

E is correct per If you run the database in NOARCHIVELOG mode, however, then the backup must be consistent; that is, you must shut down the database cleanly before the backup.

upvoted 1 times

cerebro2000x 5 years ago

CDE for me

upvoted 5 times

mik8440p 5 years, 1 month ago

I think C,E,G is correct - E - before mount you must shutdown https://docs.oracle.com/en/database/oracle/oracle-database/19/bradv/user-managed-database-backups.html#GUID-65C5E03A-E906-47EB-92AF-6DC273DBD0A8 You can make a whole database backup if a database is operating in either ARCHIVELOG or NOARCHIVELOG mode. If you run the database in NOARCHIVELOG mode, however, then the backup must be consistent; that is, you must shut down the database cleanly before the backup. "

upvoted 2 times

logica 5 years, 1 month ago

I think A B D are correct, in mount state the datafiles are closed, so you can take offline backups

upvoted 1 times

ald85 5 years, 1 month ago

ARCHIVELOG mode must be OPEN -> B is incorrect and G correct.

upvoted 1 times

veginha 5 years, 2 months ago

CDE ? C doesn't specified if the backup files where copied with the database closed. i think A and B are incorrect, the database must be closed in order to do a backup with o/s utilities

upvoted 4 times

ald85 5 years, 1 month ago

ADG A is correct, and E incorrect -> with NOARCHIVELOG must be in MOUNT state https://orahow.com/backup-oracle-database-in-noarchivelog-mode/ however with ARCHIVELOG must be OPEN, so that G is correct and B and is incorrect and D is correct too https://docs.oracle.com/html/E10643_07/rcmsynta008.htm, I think C is incorrect.

upvoted 1 times

Question #60 *Topic 1*

Examine this configuration:
1. CDB1 is a container database.
2. PDB1 and PDB2 are pluggable databases in CDB1.
You execute these commands successfully:
Which two are true? (Choose two.)

```
$ export ORACLE_SID=cdb1

$sqlplus / as sysdba

SQL> SHUTDOWN IMMEDIATE
…
Oracle instance shut down.

SQL> STARTUP MOUNT
…
Database mounted.
```

A. PDB1 and PDB2 are in MOUNT state.

B. Redo logs are opened.

C. PDB1 and PDB2 are in READ ONLY state.

D. CDB$ROOT is in MOUNT state.

E. PDB$SEED is in READ ONLY state.

Suggested Answer: *B, E*
Official Answer: BE

Community vote distribution
AD (100%)

julia Highly Voted 4 years, 7 months ago
sure is A,D:

Version 19.6.0.0.0

SQL> show pdbs

CON_ID CON_NAME OPEN MODE RESTRICTED --------- ---------- ---------- --------- 2 PDB$SEED READ ONLY YES 3 PDB1 READ WRITE YES 4 PDB2 READ WRITE YES 5 PDB3 MOUNTED SQL> shutdown immediate; Database closed. Database dismounted. ORACLE instance shut down. SQL> startup mount; ORACLE instance started. . Database mounted. SQL> show pdbs

CON_ID CON_NAME OPEN MODE RESTRICTED --------- ---------- ---------- --------- 2 PDB$SEED MOUNTED 3 PDB1 MOUNTED 4 PDB2 MOUNTED 5 PDB3 MOUNTED SQL>
upvoted 21 times

ald85 Highly Voted 4 years, 8 months ago
I think A, D are correct. When a CDB is mounted, the CBD root is mounted, which means that the control files are opened, as well as the PDBs.
upvoted 8 times

gio Most Recent 1 year, 8 months ago
Selected Answer: AD
I'm sure AD
upvoted 1 times

ScottL 1 year, 9 months ago
Selected Answer: AD

AD correct. See Julica

upvoted 1 times

jareach 2 years ago

Selected Answer: AD

In the MOUNT state, just the controlfile is opened , not redo logs nor datafiles. As the root container is mounted, i.e. not touching any datafiles, no pluggable database can be more than mounted as well.

upvoted 2 times

flaviogcmelo 2 years, 5 months ago

Selected Answer: AD

A, D are correct.

upvoted 2 times

Bio12345 2 years, 11 months ago

Selected Answer: AD

AD, demonstrated by Julica

upvoted 2 times

Dhanushbh 3 years, 6 months ago

SQL> show pdbs;

CON_ID CON_NAME OPEN MODE RESTRICTED ---------- ------------------------------ ---------- ---------- 2 PDB$SEED READ ONLY NO 3 TEST READ WRITE NO 5 TEST3 MOUNTED SQL> shutdown immediate; Database closed. Database dismounted. ORACLE instance shut down. SQL> startup mount ORACLE instance started.

Total System Global Area 6442446872 bytes Fixed Size 9154584 bytes Variable Size 3288334336 bytes Database Buffers 3137339392 bytes Redo Buffers 7618560 bytes Database mounted. SQL> show pdbs;

CON_ID CON_NAME OPEN MODE RESTRICTED ---------- ------------------------------ ---------- ---------- 2 PDB$SEED MOUNTED 3 TEST MOUNTED 5 TEST3 MOUNTED AD correct

upvoted 3 times

janw 4 years, 6 months ago

AD it is

upvoted 2 times

cerebro2000x 4 years, 6 months ago

AD for me

upvoted 3 times

Reimajestic 4 years, 7 months ago

Answer: A,D Version 19.3.0.0.0

SQL> shutdown immediate; Database closed. Database dismounted. ORACLE instance shut down.

SQL> startup mount ORACLE instance started.

Total System Global Area 2147482344 bytes Fixed Size 9136872 bytes Variable Size 637534208 bytes Database Buffers 1493172224 bytes Redo Buffers 7639040 bytes Database mounted.

SQL> SELECT NAME, CON_ID, OPEN_MODE FROM V$CONTAINERS ORDER BY CON_ID;

NAME CON_ID OPEN_MODE -------- ---------- ---------- CDB$ROOT 1 MOUNTED PDB$SEED 2 MOUNTED PDB1 3 MOUNTED PDB2 4 MOUNTED

upvoted 5 times

marcinb32 4 years, 7 months ago

AD for sure

upvoted 2 times

Question #61 *Topic 1*

Which three are true about upgrading Oracle Grid Infrastructure? (Choose three.)

A. A direct upgrade can be performed only from the immediately preceding Oracle Grid Infrastructure version.

B. The newer version is installed in a separate Oracle Grid Infrastructure home on the same server as the existing version.

C. An existing Oracle base can be used.

D. The upgrade process will automatically install all mandatory patches for the current version of Oracle Grid Infrastructure.

E. Existing Oracle Database instances must be shut down before starting the upgrade.

F. Only the grid user can perform the upgrade.

Suggested Answer: *D, E, F*
Reference: https://docs.oracle.com/database/121/CWLIN/procstop.htm#CEGEDCDB

ald85 Highly Voted 5 years, 8 months ago
I Think BCE F is incorrect, grid OS system can be named anyway. D patches won't be apply automatically A Oracle Grid only supports out-of-place upgrades
upvoted 22 times

Manojx Highly Voted 4 years, 8 months ago
BCF

https://docs.oracle.com/en/database/oracle/oracle-database/19/cwlin/restrictions-for-oracle-grid-infrastructure-upgrades.html#GUID-76797070-C2C6-4817-902C-037A67A6EF94

Oracle Grid Infrastructure upgrades are always out-of-place upgrades. You cannot perform an in-place upgrade of Oracle Grid Infrastructure to existing homes.

The same user that owned the earlier release Oracle Grid Infrastructure software must perform the Oracle Grid Infrastructure 19c upgrade.
upvoted 14 times

acesonly Most Recent 1 year, 9 months ago
Selected Answer: BCE
(Not F): "The same user that owned the earlier release Oracle Grid Infrastructure software must perform the Oracle Grid Infrastructure 19c upgrade." This formulation leaves open the possibility of having a user with a name other than 'grid'.
upvoted 1 times

gio 2 years, 8 months ago
Selected Answer: BEF
I think BEF
upvoted 1 times

vkra 3 years ago
Selected Answer: BCF
B, C, F could also correct
upvoted 1 times

vkra 2 years, 12 months ago
But I found now this, so E is really also good: https://docs.oracle.com/database/121/RILIN/procstop.htm#RILIN1173 => If you are upgrading the Oracle Clusterware or Oracle ASM software, then you must first stop all Oracle RAC database instances and all Oracle database instances that use Oracle ASM for storage. After stopping the database instances, shut down the Oracle ASM instances and Oracle Clusterware, as described in the following steps ...
upvoted 1 times

Antonio71 3 years, 1 month ago
I vote for BCE
upvoted 2 times

Cuongva 3 years, 2 months ago

Selected Answer: BCE

i think BCE

upvoted 3 times

asduta 4 years, 5 months ago

between E or F, i agree with F "The same user that owned the earlier release Oracle Grid Infrastructure software must perform the Oracle Grid Infrastructure 19c upgrade."

source : https://docs.oracle.com/en/database/oracle/oracle-database/19/cwlin/restrictions-for-oracle-grid-infrastructure-upgrades.html#GUID-76797070-C2C6-4817-902C-037A67A6EF94

upvoted 1 times

rodriguesJC 4 years, 3 months ago

E is not TRUE, because we can do a rolling upgrade

upvoted 1 times

jareach 3 years ago

Sure, but all the database instances on the currently upgraded node have to be shut down.

upvoted 2 times

JorgeZG 3 years, 10 months ago

What About Single Instances Databases??

upvoted 1 times

G_C 3 years, 4 months ago

F?? To be honest the upgrade process is doing by root user. I mean the physical process of upgrade not a prepare of home before upgrade

upvoted 1 times

asduta 3 years, 1 month ago

Yeah i agree now. Dont know why i was thinking grid owner cant be other than grid user. So F definitely wrong

upvoted 1 times

GraceYu 4 years, 7 months ago

BCD

After grid upgrade, you can list the patches:

$ORACLE_HOME/OPatch/opatch lspatches 31950577;TOO MUCH TRACING WHEN CONTENT CHECK DETECTS AND REPAIRS CORRUPTED BLOCKS 31683216;OCW Interim patch for 31683216 32240590;TOMCAT RELEASE UPDATE 19.0.0.0.0 (32240590) 32218663;ACFS RELEASE UPDATE 19.10.0.0.0 (32218663) 32218454;Database Release Update : 19.10.0.0.210119 (32218454) 29340594;DBWLM RELEASE UPDATE 19.0.0.0.0 (29340594)

upvoted 3 times

ObserverPL 4 years, 9 months ago

I agree with SimoneF - BCF.

upvoted 3 times

SimoneF 4 years, 9 months ago

I think it's BCF, cause the services are shutdown by the rootupgrade.sh:

"The script shuts down the earlier release installation, replaces it with the new Oracle Clusterware release, and starts the new Oracle Clusterware installation."

https://docs.oracle.com/en/database/oracle/oracle-database/19/cwlin/running-oui-to-upgrade-grid-infrastructure-from-an-earlier-release.html#GUID-3BE36DF3-EB62-4CD1-B53A-850A3E2129E3

While at the same time, F seems to be true since the upgrade can only be done by the same user who owned the previous installation:

"The same user that owned the earlier release Oracle Grid Infrastructure software must perform the Oracle Grid Infrastructure 19c upgrade."

https://docs.oracle.com/en/database/oracle/oracle-database/19/cwlin/restrictions-for-oracle-grid-infrastructure-upgrades.html#GUID-76797070-C2C6-4817-902C-037A67A6EF94

upvoted 6 times

SimoneF 4 years, 9 months ago

In a doc page for 12c, it's stated even more directly: "Oracle recommends that you leave Oracle RAC instances running when upgrading Oracle Clusterware. When you start the root script on each node, the database instances on that node are shut down and then the rootupgrade.sh script starts the instances again."

https://docs.oracle.com/database/121/CWLIN/procstop.htm#CWLIN519

upvoted 4 times

S4NT05 2 years, 3 months ago

F is False. You run the installer software with the Oracle grid infrastructure installation owner user account (oracle or grid).

https://docs.oracle.com/cd/E14804_01/doc/install.112/e10816/presolar.htm#CHDFAIBE

upvoted 1 times

janw 5 years, 6 months ago

bce ok

upvoted 4 times

cerebro2000x 5 years, 6 months ago

BCE for me

upvoted 2 times

julica 5 years, 7 months ago

F is incorrect because in case of 'opatchauto', root is the user that can run the command. So, I think B is right because is about new version, that means OutOfPlace Upgrade. Upgrade can reuse an existing Oracle Base (C). And E for sure, because (at the moment) GI needs DBs from that node to be down. (E). So, most probably are B,C,E

upvoted 5 times

Question #62 *Topic 1*

A database is configured in ARCHIVELOG mode.
A full RMAN backup exists but no control file backup to trace has been taken.
A media failure has occurred.
In which two scenarios is incomplete recovery required? (Choose two.)

A. after losing a SYSAUX tablespace data file

B. after losing all members of an INACTIVE online redo log group

C. after losing all members of the CURRENT online redo log group

D. after losing all copies of the control file

E. after losing an UNDO tablespace that is in use

Suggested Answer: *B, D*
Official Answer: BD

Community vote distribution
CD (75%) C (12%) CE (12%)

you1234 Highly Voted 5 years, 8 months ago
C & D is correct. C & D is correct because if you lose all control the required incomplete recovery. if incomplete recovery then db must be open in resetlog option.

E: undo tablespace. >> (Incorrect) http://oracle-help.com/backup-and-recovery/recovery-recover-loss-system-critical-undo-tablespace-original-location/

D: (correct see the details) http://oracle-help.com/backup-and-recovery/recovery-restore-loss-current-control-files-default-location/
upvoted 23 times

veginha 5 years, 8 months ago
You are right about the undo, and also the resetlog after restore the control file, but even in that case, you can do a complete recovery while you have all the redo and archives (check link) ..., I don't know which other option could be correct, so I suppose that "D" its the more accuracy.

https://web.stanford.edu/dept/itss/docs/oracle/10gR2/backup.102/b14191/rcmrecov003.htm
upvoted 1 times

janw Highly Voted 5 years, 6 months ago
C D is correct
upvoted 9 times

1451755 Most Recent 1 year, 4 months ago
Selected Answer: C
Only C is right (should be choose one). D is false, PROOF - see https://docs.oracle.com/en/database/oracle/oracle-database/19/bradv/user-managed-recovery-advanced.html#GUID-3913B513-ECF0-469A-A50A-A5A573C85DC0 Section “31.2 Recovering After the Loss of All Current Control Files” and Table “Table 31-1 Scenarios When Control Files Are Lost” - Status of Online Logs = Available, Status of Data Files = Backup, Restore Procedure = “Restore a backup control file, perform !!! COMPLETE RECOVERY !!!, and then open the database with the RESETLOGS option.” E – UNDO is also false, you can recover complete (you have anyway, full back, all archivelogs, all redo logs) see https://oracle-help.com/backup-and-recovery/recovery-recover-loss-system-critical-undo-tablespace-original-location/ Error/Mistake in Oracle University?
upvoted 1 times

gio 2 years, 8 months ago
Selected Answer: CD
https://docs.oracle.com/en/database/oracle/oracle-database/19/bradv/user-managed-recovery-advanced.html#GUID-B798CEB4-E8BE-467A-B17D-DBB993E912BE
upvoted 1 times

ScottL 2 years, 9 months ago
Selected Answer: CD

See you1234

upvoted 1 times

Antonio71 3 years, 1 month ago

C D correct

upvoted 2 times

G_C 3 years, 4 months ago

Selected Answer: CE

Teoretically find the right answers is simple but ... what if we really lost all copies of control files? Does we recover the database? In my opinion: it is not easy way but possible and it is not incomplete recovery but recreate controlfile from somewhere. But if we lost active UNDO TABLESPACE during active big transaction then we need do incomplete recovery ... I think

A. FALSE B. FALSE C. TRUE D. FALSE E. TRUE

upvoted 1 times

ErikJanssen 3 years, 3 months ago

Contents of the undo tablespace spans the transaction's DML until commit. So even if the UNDO tablespace gets lost, there is no incomplete recovery because the transaction was never committed in the first place. That was my argument for choosing answers C and D.

upvoted 2 times

mamadu 1 year, 11 months ago

D is true cause when you lose all copies of control files and assuming control file autobackup is ON, then the only way to recover the control file is to go back to a previous backup(which has the control file) and restore it then you can roll forward your database using archive logs.

upvoted 1 times

mamadu 1 year, 11 months ago

sorry no need to roll forward the database since adding this last part will make it a complete recovery.

upvoted 1 times

jareach 3 years ago

If you restore the controlfile, it is called an incomplete recovery since the database does not have the information of the last valid SCN. And you have to open the database with resetlogs.

upvoted 3 times

Bio12345 3 years, 11 months ago

Selected Answer: CD

C, D Even tough D, the redo are not lost, so the recovery would be complete. but you will need to reset the logs to open the database

upvoted 4 times

Chansi 5 years, 5 months ago

Don't think it's D as since yo have all the redo files . you just create a new control file and recover the database. Its not incomplete recovery

upvoted 3 times

jareach 3 years ago

It IS, since you do not have the information about the last valid LSN -- it was lost with the controlfiles.

upvoted 2 times

cerebro2000x 5 years, 6 months ago

CD for me

upvoted 5 times

julica 5 years, 7 months ago

(E) Undo is used for rolling BW DB. And that is useful when you want to do PITR. But in this case, UNDO does not have any role.

Incomplete recovery uses a backup to produce a noncurrent version of the database. In other words, you do not apply all of the redo records generated after the most recent backup. You usually perform incomplete recovery of the whole database in the following situations:

Media failure destroys some or all of the online redo logs. A user error causes data loss, for example, a user inadvertently drops a table. You cannot perform complete recovery because an archived redo log is missing. You lose your current control file and must use a backup control file to open the database.

In our case the qestion said there is not backup to trace for control file. So, exist a backup of CF included.

My opinion is C and D.

upvoted 5 times

taotsumiau 5 years, 7 months ago

I agreed C&D too!

upvoted 2 times

veginha 5 years, 8 months ago

I think C and E

upvoted 4 times

Question #63 *Topic 1*

Which two are true about instance recovery? (Choose two.)

A. It is not possible if an archived log is missing.

B. It is performed automatically after the database is opened; however, blocks requiring recovery are not available until they are recovered.

C. Setting FAST_START_MTTR_TARGET to a lower value reduces instance recovery time by causing dirty buffers to be written to disk more frequently, thereby reducing the number of I/Os needed during instance recovery.

D. It is performed by the Recovery Writer (RVWR) background process.

E. Setting FAST_START_MTTR_TARGET to a higher value reduces instance recovery time by causing the log writer to write more frequently, thereby reducing the number of I/Os needed during instance recovery.

F. It is performed automatically while the database remains in MOUNT state. Then the database is opened.

Suggested Answer: *E, F*

Official Answer: EF

Community vote distribution



veginha Highly Voted 6 years, 2 months ago

CF are my choices

upvoted 23 times

ald85 6 years, 2 months ago

BC Instance recovey occurs when ALTER DATABASE OPEN not before.

upvoted 9 times

veginha Highly Voted 6 years, 2 months ago

well i'm not so sure about it: You are right about the instance recovery begins after "alter database open" is executed but that doesn't mean that the database is already open when the recovery starts. The instance recovery is executed in 2 steps, rollforward and rollback and the database is open just after the rollforward is completed. So i think that is not accuracy to tell that the instance recovery is performed after the database is open. What do you thik?

Here is a good diagram to understand it. https://www.orskl.com/how-oracle-database-does-instance-recovery-after-failures/

upvoted 10 times

ald85 6 years, 2 months ago

Thanks you are correct ;)

upvoted 3 times

wagihov Most Recent 1 year, 8 months ago

Selected Answer: CF

I would exclude B as instance recovery is not related by any means to block recovery or block corruption.

upvoted 1 times

zenai 1 year, 9 months ago

Selected Answer: BC

B, C source:

upvoted 1 times

antonica 2 years, 4 months ago

CF,I Thinz

upvoted 1 times

mamadu 2 years, 5 months ago

Answer is CF:

F: instance recovery happens at mount state. from the student guide = ... The instance mounts the control files and then ATTEMPTS to open the datafiles. when it realizes the datafiles have not been synchronized during shutdown .....

B would have been correct if the word STARTED was used instead of OPENED.
upvoted 2 times

Emperor11 2 years, 8 months ago

Selected Answer: BC

Oracle Database performs instance recovery automatically in the following situations:

The database opens for the first time after the failure of a single-instance database or all instances of an Oracle RAC database. This form of instance recovery is also called crash recovery. Oracle Database recovers the online redo threads of the terminated instances together.

https://docs.oracle.com/en/database/oracle/oracle-database/19/cncpt/oracle-database-instance.html#GUID-6F20EBD4-CD7A-41BA-823B-BB3B25C7C99E

According to the documentation, B and C are true
upvoted 1 times

kaka321 3 years, 1 month ago

BC is absolutely right see below https://docs.oracle.com/cd/A57673_01/DOC/server/doc/SPS73/chap23.htm
upvoted 2 times

gio 3 years, 2 months ago

Selected Answer: CF

i'm sure of CF
upvoted 1 times

Sean82 3 years, 6 months ago

This is quite a poorly formed question in my mind, instance recovery starts when you 'startup' the database, not when it opens. If you run a shutdown abort you can clearly see in the alert log that the application of redo logs starts when the database is mounted after running the startup command and then the database is opened after it confirms that the undo tablespace is available. Then the undo is applied to rollback the database to a consistent state it was in before the unclean shutdown. So C is definitely true but B and F are both technically false as they ignore that Instance Recovery has two phases. Not sure if the Q is wrong on here but I would say that F is more true as IR starts when the DB is mounted, but it opens when rollforward is done, not when IR is completed.
upvoted 2 times

jareach 3 years, 6 months ago

Selected Answer: BC

It think it was release 8 when the faster start was introduced by opening the instance BEFORE the recovery is complete. However, the recovered blocks remain unavailable until the recovery is done.
upvoted 1 times

Antonio71 3 years, 7 months ago

For me it's BC. https://docs.oracle.com/cd/E25178_01/server.1111/e25789/startup.htm Oracle Database performs instance recovery automatically in the following situations:

The database opens for the first time after the failure of a single-instance database or all instances of an Oracle RAC database.
upvoted 1 times

Marcello86CT 3 years, 10 months ago

For me F is totally wrong. As others said the Instance recovery is performed in 2 steps , the first one is the rolling forward that needs Online Redo Logs to execute. Online Redo Logs are available just after database id open because in Mount state just control files are opened. So for me Correct answers are B & C
upvoted 4 times

kaka321 3 years, 1 month ago

I also agree with BC and your explanation
upvoted 1 times

Asaif0786 3 years, 9 months ago

if you are saying redo is applied after DB open, you got your basics very wrong.
upvoted 2 times

zenai 1 year, 9 months ago

nono, redo are applied after db is open in instance recovery. See this (the docs is for release 21, but it is the same for 19... and 12, and 11, and 10 and 9 and 8 and i'm not sure for 7 - i never covered 7) Figure 15-6 https://docs.oracle.com/en/database/oracle/oracle-database/21/cncpt/oracle-

database-instance.html#GUID-728C6BE1-5687-4DC5-B570-D2042C88F935

upvoted 1 times

Bio12345 4 years, 5 months ago

Selected Answer: CF

C F, as most agrees

upvoted 3 times

ioio 4 years, 10 months ago

BC is correct. wow. most of you are wrong. instance must be open to get undo segment online. without that, oracle doesn't know the undo data to recover the (not commited) transactions. so the phrase "remains in MOUNT state" is definitively wrong.

upvoted 5 times

niabs 5 years, 1 month ago

I will choose c and F

upvoted 4 times

ObserverPL 5 years, 3 months ago

C and F, no doubt.

upvoted 6 times

Question #64 *Topic 1*

Which two are true about Oracle Optimizer Statistics, their use, and their collection? (Choose two.)

A. The number of table rows is considered when evaluating the cost of accessing a table using an index.

B. Index balanced B*Tree height is considered when evaluating the cost of using an index.

C. The Statistics Advisor can help recommend the best way to gather statistics.

D. Statistics collected using DBMS_STATS always yield the best optimizer result.

E. The Statistics Advisor generates actions for all recommendations.

Suggested Answer: *C, D*
Reference: https://docs.oracle.com/en/database/oracle/oracle-database/19/tgsql/optimizer-statistics-advisor.html#GUID-D81A7708-FDA0-45BB-A6E2- 103858B047AE

ObserverPL Highly Voted 4 years, 3 months ago
even if A is different on the real exam, then no matters as correct answers are B and C (both according to documentation)
upvoted 12 times

veginha 5 years, 1 month ago
i believe this answer is right, CD, anyone thinks different?
upvoted 7 times

julmis Most Recent 2 years ago
answer is BC
upvoted 1 times

Guhborges 2 years, 1 month ago
Selected Answer: BC
B,C for me
upvoted 1 times

gio 2 years, 2 months ago
Selected Answer: BC
i think BC
upvoted 1 times

rodriguesJC 3 years, 9 months ago
BC for me A- Blocks not rows E- Is the same process, but the OOS are automatic and dbms_syays is manual F- "All" is so much recomendations :-)
upvoted 5 times

janw 4 years, 12 months ago
D is incorrect, proved by experience. E is incorrect: https://docs.oracle.com/en/database/oracle/oracle-database/12.2/tgsql/optimizer-statistics-advisor.html#GUID-6D8831EA-B478-4641-B84B-C341BB4D0799 When feasible, recommendations have corresponding actions.
upvoted 3 times

janw Highly Voted 4 years, 12 months ago
B is correct https://docs.oracle.com/database/121/TGSQL/tgsql_optcncpt.htm#TGSQL218 The cost of an index scan depends on the levels in the B-tree, the number of index leaf blocks to be scanned, and the number of rows to be fetched using the rowid in the index keys. The cost of fetching rows using rowids depends on the index clustering factor.
upvoted 10 times

janw 4 years, 12 months ago
A is incorrect, the number of blocks is considered instead of the number of rows. The cost of the scan depends on the number of blocks to be scanned and the multiblock read count value.
upvoted 4 times

janw 4 years, 12 months ago
C correct The output of the advisor is a report of findings and recommendations, which helps you follow best practices for gathering statistics.
upvoted 4 times

cerebro2000x 5 years ago
AD for me

upvoted 1 times

dacoben415lywenw 5 years, 1 month ago

AC obviously DBMS_STATS not always produce good statistics, especially with default options/prefs. A is not quite accurate, since cost does depend on selectivity for predicates/filters, which in turn depends on row num

upvoted 1 times

veginha 5 years, 1 month ago

option A is not in the exam. There is another response, something about the space of the table ;)

upvoted 3 times

klever35 5 years, 1 month ago

answer A D

upvoted 1 times

Question #65 *Topic 1*

A container database (CDB) contains two pluggable databases PDB1 and PDB2.
The LOCAL_UNDO_ENABLED database property is set to FALSE in the CDB.
Data file 24 of PDB2 was deleted and you need to restore and recover it.
The only RMAN backup that exists was created with the BACKUP DATABASE command while connected to CDB$ROOT.
Which three are true? (Choose three.)

A. Data file 24 can be recovered only while connected to PDB2.

B. Data file 24 can be restored and recovered while connected to CDB$ROOT.

C. Data file 24 can be restored only while connected to CDB$ROOT.

D. Data file 24 can be restored only while connected to PDB2.

E. Data file 24 can be recovered while connected to PDB2.

F. Data file 24 can be recovered while connected to CDB$ROOT.

Suggested Answer: *B, D, F*
Official Answer: BDF

Community vote distribution
BEF (100%)

ald85 Highly Voted 6 years, 2 months ago
BEF. The 'Only' word makes invalid the similar sentence.
upvoted 19 times

gio 3 years, 2 months ago
Selected Answer: BEF
I'm sure BEF. Backup database on CDB will backup all PDBs
upvoted 1 times

j_tw 3 years, 3 months ago
Selected Answer: BEF
I don't like 'only', 'always', 'must' words at this exam...
upvoted 2 times

Guhborges 3 years, 1 month ago
Me neither! B,E,F For sure
upvoted 1 times

ScottL 3 years, 3 months ago
Selected Answer: BEF
BEF correct
upvoted 1 times

jareach 3 years, 6 months ago
Selected Answer: BEF
BD are mutually exclusive.
upvoted 1 times

Manojx 3 years, 7 months ago
Per this, you must connect to CDB as target, not to PDB. Undo Mode - Shared Connect as TARGET to ... CDB root
https://docs.oracle.com/en/database/oracle/oracle-database/19/multi/backing-up-and-recovering-cdbs-and-pdbs.html#GUID-B6EF3A32-18B6-4CDC-B668-4C7B5084F371
upvoted 1 times

yumyummyyum 2 years, 11 months ago
Hi, I share your thinking, but they all said BEF. Any difference for recovering just a data file?
upvoted 1 times

7206e44 2 years, 3 months ago